

Database Questions and Answers – Relational Database and Database Schema

1. A relational database consists of a collection of

- a) Tables
- b) Fields
- c) Records
- d) Keys

[View Answer](#)

Answer: a

Explanation: Fields are the column of the relation or tables. Records are each row in a relation. Keys are the constraints in a relation.

2. A _____ in a table represents a relationship among a set of values.

- a) Column
- b) Key
- c) Row
- d) Entry

[View Answer](#)

Answer: c

Explanation: Column has only one set of values. Keys are constraints and row is one whole set of attributes. Entry is just a piece of data.

3. The term _____ is used to refer to a row.

- a) Attribute
- b) Tuple
- c) Field
- d) Instance

[View Answer](#)

Answer: b

Explanation: Tuple is one entry of the relation with several attributes which are fields.

4. The term attribute refers to a _____ of a table.

- a) Record
- b) Column
- c) Tuple
- d) Key

[View Answer](#)

Answer: b

Explanation: Attribute is a specific domain in the relation which has entries of all tuples.

5. For each attribute of a relation, there is a set of permitted values, called the _____ of that attribute.

- a) Domain
- b) Relation
- c) Set
- d) Schema

[View Answer](#)

Answer: a

Explanation: The values of the attribute should be present in the domain. Domain is a set of values permitted.

6. Database _____ which is the logical design of the database, and the database _____ which is a snapshot of the data in the database at a given instant in time.

- a) Instance, Schema

- b) Relation, Schema
- c) Relation, Domain
- d) Schema, Instance

[View Answer](#)

Answer: d

Explanation: Instance is an instance of time and schema is a representation.

7. Course(course_id,sec_id,semester)

Here the course_id,sec_id and semester are _____ and course is a _____

- a) Relations, Attribute
- b) Attributes, Relation
- c) Tuple, Relation
- d) Tuple, Attributes

[View Answer](#)

Answer: b

Explanation: The relation course has a set of attributes course_id,sec_id,semester .

8. Department (dept name, building, budget) and Employee (employee_id, name, dept name, salary)

Here the dept_name attribute appears in both the relations. Here using common attributes in relation schema is one way of relating _____ relations.

- a) Attributes of common
- b) Tuple of common
- c) Tuple of distinct
- d) Attributes of distinct

[View Answer](#)

Answer: c

Explanation: Here the relations are connected by the common attributes.

9. A domain is atomic if elements of the domain are considered to be _____ units.

- a) Different
- b) Indivisible
- c) Constant
- d) Divisible

[View Answer](#)

Answer: b

Explanation: None.

10. The tuples of the relations can be of _____ order.

- a) Any
- b) Same
- c) Sorted
- d) Constant

[View Answer](#)

Answer: a

Explanation: The values only count. The order of the tuples does not matter.

Database Questions and Answers – Keys

1. Which one of the following is a set of one or more attributes taken collectively to uniquely identify a record?

- a) Candidate key
- b) Sub key
- c) Super key
- d) Foreign key

[View Answer](#)

Answer: c

Explanation: Super key is the superset of all the keys in a relation.

2. Consider attributes ID, CITY and NAME. Which one of this can be considered as a super key?

- a) NAME
- b) ID
- c) CITY
- d) CITY, ID

[View Answer](#)

Answer: b

Explanation: Here the id is the only attribute which can be taken as a key. Other attributes are not uniquely identified.

3. The subset of a super key is a candidate key under what condition?

- a) No proper subset is a super key
- b) All subsets are super keys
- c) Subset is a super key
- d) Each subset is a super key

[View Answer](#)

Answer: a

Explanation: The subset of a set cannot be the same set. Candidate key is a set from a super key which cannot be the whole of the super set.

4. A _____ is a property of the entire relation, rather than of the individual tuples in which each tuple is unique.

- a) Rows
- b) Key
- c) Attribute
- d) Fields

[View Answer](#)

Answer: b

Explanation: Key is the constraint which specifies uniqueness.

5. Which one of the following attribute can be taken as a primary key?

- a) Name
- b) Street
- c) Id
- d) Department

[View Answer](#)

Answer: c

Explanation: The attributes name, street and department can repeat for some tuples. But the id attribute has to be unique. So it forms a primary key.

6. Which one of the following cannot be taken as a primary key?

- a) Id

- b) Register number
- c) Dept_id
- d) Street

[View Answer](#)

Answer: d

Explanation: Street is the only attribute which can occur more than once.

7. An attribute in a relation is a foreign key if the _____ key from one relation is used as an attribute in that relation.

- a) Candidate
- b) Primary
- c) Super
- d) Sub

[View Answer](#)

Answer: b

Explanation: The primary key has to be referred in the other relation to form a foreign key in that relation.

8. The relation with the attribute which is the primary key is referenced in another relation. The relation which has the attribute as a primary key is called

- a) Referential relation
- b) Referencing relation
- c) Referenced relation
- d) Referred relation

[View Answer](#)

Answer: b

Explanation: None.

9. The _____ is the one in which the primary key of one relation is used as a normal attribute in another relation.

- a) Referential relation
- b) Referencing relation
- c) Referenced relation
- d) Referred relation

[View Answer](#)

Answer: c

Explanation: None.

10. A _____ integrity constraint requires that the values appearing in specified attributes of any tuple in the referencing relation also appear in specified attributes of at least one tuple in the referenced relation.

- a) Referential
- b) Referencing
- c) Specific
- d) Primary

[View Answer](#)

Answer: a

Explanation: A relation, say r1, may include among its attributes the primary key of another relation, say r2. This attribute is called a foreign key from r1, referencing r2. The relation r1 is also called the referencing relation of the foreign key dependency, and r2 is called the referenced relation of the foreign key.

Database Questions and Answers – Relational Query Operations and Relational Operators

1. Using which language can a user request information from a database?

- a) Query
- b) Relational
- c) Structural
- d) Compiler

[View Answer](#)

Answer: a

Explanation: Query language is a method through which the database entries can be accessed.

2. Student(ID, name, dept name, tot_cred)

In this query which attributes form the primary key?

- a) Name
- b) Dept
- c) Tot_cred
- d) ID

[View Answer](#)

Answer: d

Explanation: The attributes name, dept and tot_cred can have same values unlike ID.

3. Which one of the following is a procedural language?

- a) Domain relational calculus
- b) Tuple relational calculus
- c) Relational algebra
- d) Query language

[View Answer](#)

Answer: c

Explanation: Domain and Tuple relational calculus are non-procedural language. Query language is a method through which database entries can be accessed.

4. The _____ operation allows the combining of two relations by merging pairs of tuples, one from each relation, into a single tuple.

- a) Select
- b) Join
- c) Union
- d) Intersection

[View Answer](#)

Answer: b

Explanation: Join finds the common tuple in the relations and combines it.

5. The result which operation contains all pairs of tuples from the two relations, regardless of whether their attribute values match.

- a) Join
- b) Cartesian product
- c) Intersection
- d) Set difference

[View Answer](#)

Answer: b

Explanation: Cartesian product is the multiplication of all the values in the attributes.

6. The _____ operation performs a set union of two “similarly structured” tables

- a) Union
- b) Join
- c) Product
- d) Intersect

[View Answer](#)

Answer: a

Explanation: Union just combines all the values of relations of same attributes.

7. The most commonly used operation in relational algebra for projecting a set of tuple from a relation is

- a) Join
- b) Projection
- c) Select
- d) Union

[View Answer](#)

Answer: c

Explanation: Select is used to view the tuples of the relation with or without some constraints.

8. The _____ operator takes the results of two queries and returns only rows that appear in both result sets.

- a) Union
- b) Intersect
- c) Difference
- d) Projection

[View Answer](#)

Answer: b

Explanation: The union operator gives the result which is the union of two queries and difference is the one where query which is not a part of second query.

9. A _____ is a pictorial depiction of the schema of a database that shows the relations in the database, their attributes, and primary keys and foreign keys.

- a) Schema diagram
- b) Relational algebra
- c) Database diagram
- d) Schema flow

[View Answer](#)

Answer: a

Explanation: None.

10. The _____ provides a set of operations that take one or more relations as input and return a relation as an output.

- a) Schematic representation
- b) Relational algebra
- c) Scheme diagram
- d) Relation flow

[View Answer](#)

Answer: b

Explanation: None.

Database Questions and Answers – SQL Basics and SQL Data Definition

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

1. Which one of the following is used to define the structure of the relation, deleting relations and relating schemas?

- a) DML(Data Manipulation Langauge)
- b) DDL(Data Definition Langauge)
- c) Query
- d) Relational Schema

[View Answer](#)

Answer: b

Explanation: Data Definition language is the language which performs all the operation in defining structure of relation.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

2. Which one of the following provides the ability to query information from the database and to insert tuples into, delete tuples from, and modify tuples in the database?

- a) DML(Data Manipulation Langauge)
- b) DDL(Data Definition Langauge)
- c) Query
- d) Relational Schema

[View Answer](#)

Answer: a

Explanation: DML performs the change in the values of the relation.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

What type of statement is this?

- a) DML
- b) DDL
- c) View
- d) Integrity constraint

[View Answer](#)

Answer: b

Explanation: Data Definition language is the language which performs all the operation in defining structure of relation.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

What type of statement is this?

- a) DML
- b) DDL
- c) View
- d) Integrity constraint

[View Answer](#)

Answer: a

Explanation: Select operation just shows the required fields of the relation. So it forms a DML.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

5. The basic data type char(n) is a _____ length character string and varchar(n) is _____ length character.

- a) Fixed, equal
- b) Equal, variable
- c) Fixed, variable
- d) Variable, equal

[View Answer](#)

Answer: c

Explanation: Varchar changes its length accordingly whereas char has a specific length which has to be filled by either letters or spaces.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

6. An attribute A of datatype varchar(20) has the value "Avi". The attribute B of datatype char(20) has value "Reed". Here attribute A has _____ spaces and attribute B has _____ spaces.

- a) 3, 20
- b) 20, 4
- c) 20, 20
- d) 3, 4

[View Answer](#)

Answer: a

Explanation: Varchar changes its length accordingly whereas char has a specific length which has to be filled by either letters or spaces.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

7. To remove a relation from an SQL database, we use the _____ command.

- a) Delete
- b) Purge
- c) Remove
- d) Drop table

[View Answer](#)

Answer: d

Explanation: Drop table deletes the whole structure of the relation .purge removes the table which cannot be obtained again.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

This command performs which of the following action?

- a) Remove relation
- b) Clear relation entries
- c) Delete fields
- d) Delete rows

[View Answer](#)

Answer: b

Explanation: Delete command removes the entries in the table.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

What type of statement is this?

- a) Query
- b) DML
- c) Relational
- d) DDL

[View Answer](#)

Answer: b

Explanation: The values are manipulated. So it is a DML.

3.

```
CREATE TABLE employee (name VARCHAR, id INTEGER)
```

10. Updates that violate _____ are disallowed.

- a) Integrity constraints
- b) Transaction control
- c) Authorization
- d) DDL constraints

[View Answer](#)

Answer: a

Explanation: Integrity constraint has to be maintained in the entries of the relation.

ayzom.com

Database Questions and Answers – SQL Queries

Which of these query will display the the table given above ?

- a) Select employee from name
- b) Select name
- c) Select name from employee
- d) Select employee

[View Answer](#)

Answer: c

Explanation: The field to be displayed is included in select and the table is included in the from clause.

- a) All
- b) From
- c) Distinct
- d) Name

[View Answer](#)

Answer: c

Explanation: Distinct keyword selects only the entries that are unique.

3. The _____ clause allows us to select only those rows in the result relation of the _____ clause that satisfy a specified predicate.

- a) Where, from
- b) From, select
- c) Select, from
- d) From, where

[View Answer](#)

Answer: a

Explanation: Where selects the rows on a particular condition. From gives the relation which involves the operation.

- a) Salary*1.1
- b) ID
- c) Where
- d) Instructor

[View Answer](#)

Answer: c

Explanation: Where selects the rows on a particular condition. From gives the relation which involves the operation. Since Instructor is a relation it has to have from clause.

5. The _____ clause is used to list the attributes desired in the result of a query.

- a) Where
- b) Select
- c) From
- d) Distinct

[View Answer](#)

Answer: b

Explanation: None

- a) Select name, course_id from teaches, instructor where instructor_id=course_id;
- b) Select name, course_id from instructor natural join teaches;
- c) Select name, course_id from instructor;
- d) Select course_id from instructor join teaches;

[View Answer](#)

Answer: b

Explanation: Join clause joins two tables by matching the common column.

Which of the following fields are displayed as output?

- a) Salary, dept_id
- b) Employee
- c) Salary
- d) All the field of employee relation

[View Answer](#)

Answer: d

Explanation: Here * is used to select all the fields of the relation.

- a) 1009, 1001, 1018
- b) 1009, 1018
- c) 1001
- d) 1018

[View Answer](#)

Answer: d

Explanation: Greater than symbol does not include the given value unlike >=.

9. Which of the following statements contains an error?

- a) Select * from emp where empid = 10003;
- b) Select empid from emp where empid = 10006;
- c) Select empid from emp;
- d) Select empid where empid = 1009 and lastname = 'GELLER';

[View Answer](#)

Answer: d

Explanation: This query do not have from clause which specifies the relation from which the values has to be selected.

- a) Table
- b) Values
- c) Relation
- d) Field

[View Answer](#)

Answer: b

Explanation: Value keyword has to be used to insert the values into the table.

Database Questions and Answers – Basic SQL Operations

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE ' _____ Computer Science';
```

Which keyword must be used here to rename the field name?

- a) From
- b) Rename
- c) As
- d) Join

[View Answer](#)

Answer: c

Explanation: As keyword is used to rename.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE ' _____ Computer Science';
```

In the SQL given above there is an error . Identify the error.

- a) Dept_name
- b) Employee
- c) "Comp Sci"
- d) From

[View Answer](#)

Answer: c

Explanation: For any string operations single quoted(') must be used to enclose.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE ' _____ Computer Science';
```

Which one of the following has to be added into the blank to select the dept_name which has Computer Science as its ending string?

- a) %
- b) _
- c) ||
- d) \$

[View Answer](#)

Answer: a

Explanation: The % character matches any substring.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE ' _____ Computer Science';
```

4. ' ___ ' matches any string of _____ three characters. ' ___ %' matches any string of at _____ three characters.

- a) Atleast, Exactly
- b) Exactly, Atleast

- c) Atleast, All
- d) All, Exactly

[View Answer](#)

Answer: b

Explanation: None.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE '_____ Computer Science';
```

By default, the order by clause lists items in _____ order.

- a) Descending
- b) Any
- c) Same
- d) Ascending

[View Answer](#)

Answer: d

Explanation: Specification of descending order is essential but it not for ascending.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE '_____ Computer Science';
```

To display the salary from greater to smaller and name in ascending order which of the following options should be used?

- a) Ascending, Descending
- b) Asc, Desc
- c) Desc, Asc
- d) Descending, Ascending

[View Answer](#)

Answer: c

Explanation: None.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE '_____ Computer Science';
```

[View Answer](#)

Answer: a

Explanation: SQL includes a between comparison operator to simplify where clauses that specify that a value be less than or equal to some value and greater than or equal to some other value.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE '_____ Computer Science';
```

This query does which of the following operation?

- a) All attributes of instructor and teaches are selected
- b) All attributes of instructor are selected on the given condition
- c) All attributes of teaches are selected on given condition
- d) Only the some attributes from instructed and teaches are selected

[View Answer](#)

Answer: b

Explanation: The asterisk symbol “ * ” can be used in the select clause to denote “all attributes.”

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE ' _____ Computer Science';
```

9. In SQL the spaces at the end of the string are removed by _____ function.

- a) Upper
- b) String
- c) Trim
- d) Lower

[View Answer](#)

Answer: c

Explanation: The syntax of trim is Trim(s); where s-string.

3.

```
SELECT emp_name
FROM department
WHERE dept_name LIKE ' _____ Computer Science';
```

10. _____ operator is used for appending two strings.

- a) &
- b) %
- c) ||
- d) _

[View Answer](#)

Answer: c

Explanation: || is the concatenation operator.

Database Questions and Answers – Set Operations

1. The union operation is represented by

- a) \cap
- b) U
- c) $-$
- d) $*$

[View Answer](#)

Answer: b

Explanation: Union operator combines the relations.

2. The intersection operator is used to get the _____ tuples.

- a) Different
- b) Common
- c) All
- d) Repeating

[View Answer](#)

Answer: b

Explanation: Intersection operator ignores unique tuples and takes only common ones.

3. The union operation automatically _____ unlike the select clause.

- a) Adds tuples
- b) Eliminates unique tuples
- c) Adds common tuples
- d) Eliminates duplicate

[View Answer](#)

Answer: d

Explanation: None.

4. If we want to retain all duplicates, we must write _____ in place of union.

- a) Union all
- b) Union some
- c) Intersect all
- d) Intersect some

[View Answer](#)

Answer: a

Explanation: Union all will combine all the tuples including duplicates.

This query displays

- a) Only tuples from second part
- b) Only tuples from the first part which has the tuples from second part
- c) Tuples from both the parts
- d) Tuples from first part which do not have second part

[View Answer](#)

Answer: d

Explanation: Except keyword is used to ignore the values.

- a) i-only
- b) ii-only
- c) i & ii

d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: % is used with like and _ is used to fill in the character.

7. The number of attributes in relation is called as its

- a) Cardinality
- b) Degree
- c) Tuples
- d) Entity

[View Answer](#)

Answer: b

Explanation: None.

8. _____ clause is an additional filter that is applied to the result.

- a) Select
- b) Group-by
- c) Having
- d) Order by

[View Answer](#)

Answer: c

Explanation: Having is used to provide additional aggregate filtration to the query.

9. _____ joins are SQL server default

- a) Outer
- b) Inner
- c) Equi
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: It is optional to give the inner keyword with the join as it is default.

10. The _____ is essentially used to search for patterns in target string.

- a) Like Predicate
- b) Null Predicate
- c) In Predicate
- d) Out Predicate

[View Answer](#)

Answer: a

Explanation: Like predicate matches the string in the given pattern.

Database Questions and Answers – Null Values Operations

1. A _____ indicates an absent value that may exist but be unknown or that may not exist at all.

- a) Empty tuple
- b) New value
- c) Null value
- d) Old value

[View Answer](#)

Answer: c

Explanation: None.

2. If the attribute phone number is included in the relation all the values need not be entered into the phone number column. This type of entry is given as

- a) 0
- b) –
- c) Null
- d) Empty space

[View Answer](#)

Answer: c

Explanation: Null is used to represent the absence of a value.

3. The predicate in a where clause can involve Boolean operations such as and. The result of true and unknown is _____ false and unknown is _____ while unknown and unknown is _____

- a) Unknown, unknown, false
- b) True, false, unknown
- c) True, unknown, unknown
- d) Unknown, false, unknown

[View Answer](#)

Answer: d

Explanation: None.

- a) Tuples with null value
- b) Tuples with no null values
- c) Tuples with any salary
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: Not null constraint removes the tuples of null values.

5. In an employee table to include the attributes whose value always have some value which of the following constraint must be used?

- a) Null
- b) Not null
- c) Unique
- d) Distinct

[View Answer](#)

Answer: b

Explanation: Not null constraint removes the tuples of null values.

6. Using the _____ clause retains only one copy of such identical tuples.

- a) Null

- b) Unique
- c) Not null
- d) Distinct

[View Answer](#)

Answer: d

Explanation: Unique is a constraint.

Some of these insert statements will produce an error. Identify the statement.

- a) Insert into employee values (1005,Rach,0);
- b) Insert into employee values (1002,Joey,335);
- c) Insert into employee values (1007,Ross,);
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Not null constraint is specified which means some value (can include 0 also) should be given.

8. The primary key must be

- a) Unique
- b) Not null
- c) Both Unique and Not null
- d) Either Unique or Not null

[View Answer](#)

Answer: c

Explanation: Primary key must satisfy unique and not null condition for sure.

Why does this statement cause an error when QUANTITY values are null?

- a) The expression attempts to divide by a null value
- b) The data types in the conversion function are incompatible
- c) The character string none should be enclosed in single quotes (' ')
- d) A null value used in an expression cannot be converted to an actual value

[View Answer](#)

Answer: a

Explanation: The expression attempts to divide by a null value is erroneous in sql.

10. The result of _____ unknown is unknown.

- a) Xor
- b) Or
- c) And
- d) Not

[View Answer](#)

Answer: d

Explanation: Since unknown does not hold any value the value cannot have a reverse value.

Database Questions and Answers – Aggregate Functions and Nested Subqueries – 1

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

1. Aggregate functions are functions that take a _____ as input and return a single value.

- a) Collection of values
- b) Single value
- c) Aggregate value
- d) Both Collection of values & Single value

[View Answer](#)

Answer: a

Explanation: None.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

Which of the following should be used to find the mean of the salary ?

- a) Mean(salary)
- b) Avg(salary)
- c) Sum(salary)
- d) Count(salary)

[View Answer](#)

Answer: b

Explanation: Avg() is used to find the mean of the values.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

If we do want to eliminate duplicates, we use the keyword _____ in the aggregate expression.

- a) Distinct
- b) Count
- c) Avg
- d) Primary key

[View Answer](#)

Answer: a

Explanation: Distinct keyword is used to select only unique items from the relation.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

4. All aggregate functions except _____ ignore null values in their input collection.

- a) Count(attribute)
- b) Count(*)

c) Avg

d) Sum

[View Answer](#)

Answer: b

Explanation: * is used to select all values including null.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

5. A Boolean data type that can take values true, false, and_____

a) 1

b) 0

c) Null

d) Unknown

[View Answer](#)

Answer: d

Explanation: Unknown values do not take null value but it is not known.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

6. The ____ connective tests for set membership, where the set is a collection of values produced by a select clause. The ____ connective tests for the absence of set membership.

a) Or, in

b) Not in, in

c) In, not in

d) In, or

[View Answer](#)

Answer: c

Explanation: In checks, if the query has the value but not in checks if it does not have the value.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

[View Answer](#)

Answer: a

Explanation: None.

3.

```
SELECT COUNT (____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

8. The phrase “greater than at least one” is represented in SQL by_____

a) < all

b) < some

c) > all

d) > some

[View Answer](#)

Answer: d

Explanation: >some takes atleast one value above it .

3.

```
SELECT COUNT (_____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

[View Answer](#)

Answer: a

Explanation: None.

3.

```
SELECT COUNT (_____ ID)
FROM teaches
WHERE semester = 'Spring' AND YEAR = 2010;
```

10. We can test for the nonexistence of tuples in a subquery by using the _____ construct.

- a) Not exist
- b) Not exists
- c) Exists
- d) Exist

[View Answer](#)

Answer: b

Explanation: Exists is used to check for the existence of tuples.

Database Questions and Answers – Aggregate Functions and Nested Subqueries – 2

- a) Avg(salary) should not be selected
- b) Dept_id should not be used in group by clause
- c) Misplaced group by clause
- d) Group by clause is not valid in this query

[View Answer](#)

Answer: b

Explanation: Any attribute that is not present in the group by clause must appear only inside an aggregate function if it appears in the select clause, otherwise the query is treated as erroneous.

2. SQL applies predicates in the _____ clause after groups have been formed, so aggregate functions may be used.

- a) Group by
- b) With
- c) Where
- d) Having

[View Answer](#)

Answer: b

Explanation: The with clause provides away of defining a temporary relation whose definition is available only to the query in which the with clause occurs.

3. Aggregate functions can be used in the select list or the _____ clause of a select statement or subquery. They cannot be used in a _____ clause.

- a) Where, having
- b) Having, where
- c) Group by, having
- d) Group by, where

[View Answer](#)

Answer: b

Explanation: To include aggregate functions having clause must be included after where.

4. The _____ keyword is used to access attributes of preceding tables or subqueries in the from clause.

- a) In
- b) Lateral
- c) Having
- d) With

[View Answer](#)

Answer: b

Explanation:

```
Eg : SELECT name, salary, avg salary
      FROM instructor I1, lateral (SELECT avg(salary) AS avg salary
      FROM instructor I2
      WHERE I2.dept name= I1.dept name);
```

Without the lateral clause, the subquery cannot access the correlation variable I1 from the outer query.

5. Which of the following creates a temporary relation for the query on which it is defined?

- a) With
- b) From
- c) Where

d) Select

[View Answer](#)

Answer: a

Explanation: The with clause provides a way of defining a temporary relation whose definition is available only to the query in which the with clause occurs.

In the query given above which one of the following is a temporary relation?

- a) Budget
- b) Department
- c) Value
- d) Max_budget

[View Answer](#)

Answer: d

Explanation: With clause creates a temporary relation.

7. Subqueries cannot:

- a) Use group by or group functions
- b) Retrieve data from a table different from the one in the outer query
- c) Join tables
- d) Appear in select, update, delete, insert statements.

[View Answer](#)

Answer: c

Explanation: None.

8. Which of the following is not an aggregate function?

- a) Avg
- b) Sum
- c) With
- d) Min

[View Answer](#)

Answer: c

Explanation: With is used to create temporary relation and its not an aggregate function.

9. The EXISTS keyword will be true if:

- a) Any row in the subquery meets the condition only
- b) All rows in the subquery fail the condition only
- c) Both of these two conditions are met
- d) Neither of these two conditions is met

[View Answer](#)

Answer: a

Explanation: EXISTS keyword checks for existance of a condition.

10. How can you find rows that do not match some specified condition?

- a) EXISTS
- b) Double use of NOT EXISTS
- c) NOT EXISTS
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: None.

Database Questions and Answers – Modification of Database

1. A Delete command operates on _____ relation.

- a) One
- b) Two
- c) Several
- d) Null

[View Answer](#)

Answer: a

Explanation: Delete can delete from only one table at a time.

The above command

- a) Deletes a particular tuple from the relation
- b) Deletes the relation
- c) Clears all entries from the relation
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Here P gives the condition for deleting specific rows.

3. Which one of the following deletes all the entries but keeps the structure of the relation.

- a) Delete from r where P;
- b) Delete from instructor where dept name= 'Finance';
- c) Delete from instructor where salary between 13000 and 15000;
- d) Delete from instructor;

[View Answer](#)

Answer: d

Explanation: Absence of condition deletes all rows.

d) Not possible

[View Answer](#)

Answer: b

Explanation: Using select statement in insert will include rows which are the result of the selection.

d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: The query must include building=watson condition to filter the tuples.

Fill in with correct keyword to update the instructor relation.

- a) Where
- b) Set
- c) In
- d) Select

[View Answer](#)

Answer: b

Explanation: Set is used to update the particular value.

7. _____ are useful in SQL update statements, where they can be used in the set clause.

- a) Multiple queries
- b) Sub queries
- c) Update
- d) Scalar subqueries

[View Answer](#)

Answer: d

Explanation: None.

8. The problem of ordering the update in multiple updates is avoided using

- a) Set
- b) Where
- c) Case
- d) When

[View Answer](#)

Answer: c

Explanation: The case statements can add the order of updating tuples.

- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: None.

- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: The order of the two update statements is important. If we changed the order of the two statements, an instructor with a salary just under \$100,000 would receive an over 8 percent raise. SQL provides a case construct that we can use to perform both the updates with a single update statement, avoiding the problem with the order of updates.

Database Questions and Answers – Join Expressions

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

1. The _____ condition allows a general predicate over the relations being joined.

- a) On
- b) Using
- c) Set
- d) Where

[View Answer](#)

Answer: a

Explanation: On gives the condition for the join expression.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

2. Which of the join operations do not preserve non matched tuples?

- a) Left outer join
- b) Right outer join
- c) Inner join
- d) Natural join

[View Answer](#)

Answer: c

Explanation: INNER JOIN: Returns all rows when there is at least one match in BOTH tables.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Join can be replaced by inner join.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

4. What type of join is needed when you wish to include rows that do not have matching values?

- a) Equi-join
- b) Natural join
- c) Outer join
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: An outer join does not require each record in the two joined tables to have a matching record..

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

5. How many tables may be included with a join?

- a) One
- b) Two
- c) Three
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: Join can combine multiple tables.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

6. Which are the join types in join condition:

- a) Cross join
- b) Natural join
- c) Join with USING clause
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: There are totally four join types in SQL.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

7. How many join types in join condition:

- a) 2
- b) 3
- c) 4
- d) 5

[View Answer](#)

Answer: d

Explanation: Types are inner join, left outer join, right outer join, full join, cross join.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

8. Which join refers to join records from the right table that have no matching key in the left table are include in the result set:

- a) Left outer join
- b) Right outer join
- c) Full outer join
- d) Half outer join

[View Answer](#)

Answer: b

Explanation: RIGHT OUTER JOIN: Return all rows from the right table and the matched rows from the left table.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

9. The operation which is not considered a basic operation of relational algebra is

- a) Join
- b) Selection
- c) Union
- d) Cross product

[View Answer](#)

Answer: a

Explanation: None.

c)

```
SELECT *
FROM student LEFT OUTER JOIN takes USING (ID);
```

10. In SQL the statement select * from R, S is equivalent to

- a) Select * from R natural join S
- b) Select * from R cross join S
- c) Select * from R union join S
- d) Select * from R inner join S

[View Answer](#)

Answer: b

Explanation: None.

Database Questions and Answers – Views

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

1. Which of the following creates a virtual relation for storing the query?

- a) Function
- b) View
- c) Procedure
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Any such relation that is not part of the logical model, but is made visible to a user as a virtual relation, is called a view.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

2. Which of the following is the syntax for views where v is view name?

- a) Create view v as “query name”;
- b) Create “query expression” as view;
- c) Create view v as “query expression”;
- d) Create view “query expression”;

[View Answer](#)

Answer: c

Explanation: <query expression> is any legal query expression. The view name is represented by v.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

Here the tuples are selected from the view. Which one denotes the view.

- a) Course_id
- b) Watson
- c) Building
- d) physics_fall_2009

[View Answer](#)

Answer: c

Explanation: View names may appear in a query any place where a relation name may appear.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

4. Materialised views make sure that

- a) View definition is kept stable
- b) View definition is kept up-to-date

- c) View definition is verified for error
- d) View is deleted after specified time

[View Answer](#)

Answer: b

Explanation: None.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

5. Updating the value of the view

- a) Will affect the relation from which it is defined
- b) Will not change the view definition
- c) Will not affect the relation from which it is defined
- d) Cannot determine

[View Answer](#)

Answer: a

Explanation: None.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

6. SQL view is said to be updatable (that is, inserts, updates or deletes can be applied on the view) if which of the following conditions are satisfied by the query defining the view?

- a) The from clause has only one database relation
- b) The query does not have a group by or having clause
- c) The select clause contains only attribute names of the relation and does not have any expressions, aggregates, or distinct specification
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: All of the conditions must be satisfied to update the view in sql.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

7. Which of the following is used at the end of the view to reject the tuples which do not satisfy the condition in where clause?

- a) With
- b) Check
- c) With check
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Views can be defined with a with check option clause at the end of the view definition; then, if a tuple inserted into the view does not satisfy the view's where clause condition, the insertion is rejected by the database system.

3.

```
SELECT course_id
```

```
FROM physics_fall_2009
WHERE building= 'Watson';
```

[View Answer](#)

Answer: a

Explanation: None.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

If we insert tuple into the view as insert into instructor info values ('69987', 'White', 'Taylor');
What will be the values of the other attributes in instructor and department relations?

- a) Default value
- b) Null
- c) Error statement
- d) 0

[View Answer](#)

Answer: b

Explanation: The values take null if there is no constraint in the attribute else it is an Erroneous statement.

3.

```
SELECT course_id
FROM physics_fall_2009
WHERE building= 'Watson';
```

Find the error in this query.

- a) Instructor
- b) Select
- c) View ...as
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Syntax is – create view v as <query expression>;

Database Questions and Answers – Transactions

1. A _____ consists of a sequence of query and/or update statements.

- a) Transaction
- b) Commit
- c) Rollback
- d) Flashback

[View Answer](#)

Answer: a

Explanation: Transaction is a set of operation until commit.

2. Which of the following makes the transaction permanent in the database?

- a) View
- b) Commit
- c) Rollback
- d) Flashback

[View Answer](#)

Answer: b

Explanation: Commit work commits the current transaction.

3. In order to undo the work of transaction after last commit which one should be used?

- a) View
- b) Commit
- c) Rollback
- d) Flashback

[View Answer](#)

Answer: c

Explanation: Rollback work causes the current transaction to be rolled back; that is, it undoes all the updates performed by the SQL statements in the transaction.

What does Rollback do?

- a) Undoes the transactions before commit
- b) Clears all transactions
- c) Redoes the transactions before commit
- d) No action

[View Answer](#)

Answer: d

Explanation: Once a transaction has executed commit work, its effects can no longer be undone by rollback work.

5. In case of any shut down during transaction before commit which of the following statement is done automatically?

- a) View
- b) Commit
- c) Rollback
- d) Flashback

[View Answer](#)

Answer: c

Explanation: Once a transaction has executed commit work, its effects can no longer be undone by rollback work.

6. In order to maintain the consistency during transactions, database provides

- a) Commit

- b) Atomic
- c) Flashback
- d) Retain

[View Answer](#)

Answer: b

Explanation: By atomic, either all the effects of the transaction are reflected in the database, or none are (after rollback).

7. Transaction processing is associated with everything below except

- a) Conforming an action or triggering a response
- b) Producing detail summary or exception report
- c) Recording a business activity
- d) Maintaining a data

[View Answer](#)

Answer: a

Explanation: None.

8. A transaction completes its execution is said to be

- a) Committed
- b) Aborted
- c) Rolled back
- d) Failed

[View Answer](#)

Answer: a

Explanation: A complete transaction always commits.

9. Which of the following is used to get back all the transactions back after rollback?

- a) Commit
- b) Rollback
- c) Flashback
- d) Redo

[View Answer](#)

Answer: c

Explanation: None.

10. _____ will undo all statements up to commit?

- a) Transaction
- b) Flashback
- c) Rollback
- d) Abort

[View Answer](#)

Answer: c

Explanation: Flashback will undo all the statements and Abort will terminate the operation.

Database Questions and Answers – Integrity Constraints

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

1. To include integrity constraint in an existing relation use :

- a) Create table
- b) Modify table
- c) Alter table
- d) Drop table

[View Answer](#)

Answer: c

Explanation: SYNTAX – alter table table-name add constraint, where constraint can be any constraint on the relation.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

2. Which of the following is not an integrity constraint?

- a) Not null
- b) Positive
- c) Unique
- d) Check ‘predicate’

[View Answer](#)

Answer: b

Explanation: Positive is a value and not a constraint.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

What will be the result of the query?

- a) All statements executed
- b) Error in create statement
- c) Error in insert into Employee values(1006,Ted,Finance,);
- d) Error in insert into Employee values(1008,Ross,Sales,20000);

[View Answer](#)

Answer: d

Explanation: The not null specification prohibits the insertion of a null value for the attribute.

The unique specification says that no two tuples in the relation can be equal on all the listed attributes.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

Inorder to ensure that the value of budget is non-negative which of the following should be used?

- a) Check(budget>0)
- b) Check(budget<0)
- c) Alter(budget>0)
- d) Alter(budget<0)

[View Answer](#)

Answer: a

Explanation: A common use of the check clause is to ensure that attribute values satisfy specified conditions, in effect creating a powerful type system.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

5. Foreign key is the one in which the _____ of one relation is referenced in another relation.

- a) Foreign key
- b) Primary key
- c) References
- d) Check constraint

[View Answer](#)

Answer: b

Explanation: The foreign-key declaration specifies that for each course tuple, the department name specified in the tuple must exist in the department relation.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

Which of the following is used to delete the entries in the referenced table when the tuple is deleted in course table?

- a) Delete
- b) Delete cascade
- c) Set null
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: The delete “cascades” to the course relation, deletes the tuple that refers to the department that was deleted.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

7. Domain constraints, functional dependency and referential integrity are special forms of _____

- a) Foreign key
- b) Primary key
- c) Assertion
- d) Referential constraint

[View Answer](#)

Answer: c

Explanation: An assertion is a predicate expressing a condition we wish the database to always satisfy.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

8. Which of the following is the right syntax for the assertion?

- a) Create assertion 'assertion-name' check 'predicate';
- b) Create assertion check 'predicate' 'assertion-name';
- c) Create assertions 'predicates';
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: None.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

9. Data integrity constraints are used to:

- a) Control who is allowed access to the data
- b) Ensure that duplicate records are not entered into the table
- c) Improve the quality of data entered for a specific property (i.e., table column)
- d) Prevent users from changing the values stored in the table

[View Answer](#)

Answer: c

Explanation: None.

3.

```
CREATE TABLE Employee(Emp_id NUMERIC NOT NULL, Name VARCHAR(20) , dept_name VARCHAR(20), Salary NUMERIC UNIQUE
INSERT INTO Employee VALUES(1002, Ross, CSE, 10000)
INSERT INTO Employee VALUES(1006,Ted,Finance, );
INSERT INTO Employee VALUES(1002,Rita,Sales,20000);
```

10. Which of the following can be addressed by enforcing a referential integrity constraint?

- a) All phone numbers must include the area code
- b) Certain fields are required (such as the email address, or phone number) before the record is accepted
- c) Information on the customer must be known before anything can be sold to that customer
- d) When entering an order quantity, the user must input a number and not some text (i.e., 12 rather than 'a dozen')

[View Answer](#)

Answer: c

Explanation: The information can be referred to and obtained.

Database Questions and Answers – SQL Data Types and Schemas

1. Dates must be specified in the format

- a) mm/dd/yy
- b) yyyy/mm/dd
- c) dd/mm/yy
- d) yy/dd/mm

[View Answer](#)

Answer: b

Explanation: yyyy/mm/dd is the default format in sql.

2. A _____ on an attribute of a relation is a data structure that allows the database system to find those tuples in the relation that have a specified value for that attribute efficiently, without scanning through all the tuples of the relation.

- a) Index
- b) Reference
- c) Assertion
- d) Timestamp

[View Answer](#)

Answer: a

Explanation: Index is the reference to the tuples in a relation.

Here which one denotes the relation for which index is created?

- a) StudentID_index
- b) ID
- c) StudentID
- d) Student

[View Answer](#)

Answer: d

Explanation: The statement creates an index named studentID index on the attribute ID of the relation student.

4. Which of the following is used to store movie and image files?

- a) Clob
- b) Blob
- c) Binary
- d) Image

[View Answer](#)

Answer: b

Explanation: SQL therefore provides large-object data types for character data (clob) and binary data (blob). The letters “lob” in these data types stand for “Large OBject”.

5. The user defined data type can be created using

- a) Create datatype
- b) Create data
- c) Create definetype
- d) Create type

[View Answer](#)

Answer: d

Explanation: The create type clause can be used to define new types. Syntax : create type Dollars as numeric(12,2) final; .

6. Values of one type can be converted to another domain using which of the following?

- a) Cast
- b) Drop type
- c) Alter type
- d) Convert

[View Answer](#)

Answer: a

Explanation: Example of cast :cast (department.budget to numeric(12,2)). SQL provides drop type and alter type clauses to drop or modify types that have been created earlier.

In order to ensure that an instructor's salary domain allows only values greater than a specified value use:

- a) Value>=30000.00
- b) Not null;
- c) Check(value >= 29000.00);
- d) Check(value)

[View Answer](#)

Answer: c

Explanation: Check(value 'condition') is the syntax.

8. Which of the following closely resembles Create view?

- a) Create table . . . like
- b) Create table . . . as
- c) With data
- d) Create view as

[View Answer](#)

Answer: b

Explanation: The 'create table . . . as' statement closely resembles the create view statement and both are defined by using queries. The main difference is that the contents of the table are set when the table is created, whereas the contents of a view always reflect the current query result.

9. In contemporary databases, the top level of the hierarchy consists of _____ each of which can contain _____

- a) Catalogs, schemas
- b) Schemas, catalogs
- c) Environment, schemas
- d) Schemas, Environment

[View Answer](#)

Answer: a

Explanation: None.

10. Which of the following statements creates a new table temp_instructor that has the same schema as an instructor.

- a) create table temp_instructor;
- b) Create table temp_instructor like instructor;
- c) Create Table as temp_instructor;
- d) Create table like temp_instructor;

[View Answer](#)

Answer: b

Explanation: None.

Database Questions and Answers – Authorizations

1. The database administrator who authorizes all the new users, modifies the database and takes grants privilege is

- a) Super user
- b) Administrator
- c) Operator of operating system
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: The authorizations provided by the administrator to the user is a privilege.

[View Answer](#)

Answer: a

Explanation: The privilege list allows the granting of several privileges in one command .

3. Which of the following is used to provide privilege to only a particular attribute?

- a) Grant select on employee to Amit
- b) Grant update(budget) on department to Raj
- c) Grant update(budget,salary,Rate) on department to Raj
- d) Grant delete to Amit

[View Answer](#)

Answer: b

Explanation: This grant statement gives user Raj update authorization on the budget attribute of the department relation.

4. Which of the following statement is used to remove the privilege from the user Amir?

- a) Remove update on department from Amir
- b) Revoke update on employee from Amir
- c) Delete select on department from Raj
- d) Grant update on employee from Amir

[View Answer](#)

Answer: b

Explanation: revoke on from ;

d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: The role is first created and the authorization is given on relation takes to the role.

6. Which of the following is true regarding views?

- a) The user who creates a view cannot be given update authorization on a view without having update authorization on the relations used to define the view
- b) The user who creates a view cannot be given update authorization on a view without having update authorization on the relations used to define the view
- c) If a user creates a view on which no authorization can be granted, the system will allow the view creation request
- d) A user who creates a view receives all privileges on that view

[View Answer](#)

Answer: c

Explanation: A user who creates a view does not necessarily receive all privileges on that view.

7. If we wish to grant a privilege and to allow the recipient to pass the privilege on to other users, we append the _____ clause to the

appropriate grant command.

- a) With grant
- b) Grant user
- c) Grant pass privilege
- d) With grant option

[View Answer](#)

Answer: d

Explanation: None.

8. In authorization graph, if DBA provides authorization to u1 which inturn gives to u2 which of the following is correct?

- a) If DBA revokes authorization from u1 then u2 authorization is also revoked
- b) If u1 revokes authorization from u2 then u2 authorization is revoked
- c) If DBA & u1 revokes authorization from u1 then u2 authorization is also revoked
- d) If u2 revokes authorization then u1 authorization is revoked

[View Answer](#)

Answer: c

Explanation: A user has an authorization if and only if there is a path from the root of the authorization graph down to the node representing the user.

9. Which of the following is used to avoid cascading of authorizations from the user?

- a) Granted by current role
- b) Revoke select on department from Amit, Satoshi restrict;
- c) Revoke grant option for select on department from Amit;
- d) Revoke select on department from Amit, Satoshi cascade;

[View Answer](#)

Answer: b

Explanation: The revoke statement may specify restrict in order to prevent cascading revocation. The keyword cascade can be used instead of restrict to indicate that revocation should cascade.

10. The granting and revoking of roles by the user may cause some confusions when that user role is revoked. To overcome the above situation

- a) The privilege must be granted only by roles
- b) The privilege is granted by roles and users
- c) The user role cannot be removed once given
- d) By restricting the user access to the roles

[View Answer](#)

Answer: a

Explanation: The current role associated with a session can be set by executing set role name. The specified role must have been granted to the user, else the set role statement fails.

Database Questions and Answers – Access SQL From a Programming Language

1. Which of the following is used to access the database server at the time of executing the program and get the data from the server accordingly?

- a) Embedded SQL
- b) Dynamic SQL
- c) SQL declarations
- d) SQL data analysis

[View Answer](#)

Answer: b

Explanation: Embedded SQL, the SQL statements are identified at compile time using a preprocessor. The preprocessor submits the SQL statements to the database system for precompilation and optimization; then it replaces the SQL statements in the application program with appropriate code and function calls before invoking the programming-language compiler.

2. Which of the following header must be included in java program to establish database connectivity using JDBC ?

- a) Import java.sql.*;
- b) Import java.sql.jdbc.jdbc.*;
- c) Import java.jdbc.*;
- d) Import java.sql.jdbc.*;

[View Answer](#)

Answer: a

Explanation: The Java program must import `java.sql.*`, which contains the interface definitions for the functionality provided by JDBC.

3. DriverManager.getConnection(_____, _____, _____)

What are the two parameters that are included?

- a) URL or machine name where server runs, Password, User ID
- b) URL or machine name where server runs, User ID, Password
- c) User ID, Password, URL or machine name where server runs
- d) Password, URL or machine name where server runs, User ID

[View Answer](#)

Answer: b

Explanation: The database must be opened first in order to perform any operations for which this get connection method is used.

4. Which of the following invokes functions in sql?

- a) Prepared Statements
- b) Connection statement
- c) Callable statements
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: JDBC provides a Callable Statement interface that allows invocation of SQL stored procedures and functions.

5. Which of the following function is used to find the column count of the particular resultset?

- a) getMetaData()
- b) Metadata()
- c) getColumn()
- d) get Count()

[View Answer](#)

Answer: a

Explanation: The interface ResultSet

has a method, `getMetaData()`, that returns a `ResultSetMetaData` object that contains metadata about the result set. `ResultSetMetaData`, in turn, has methods to find metadata information, such as the number of columns in the result, the name of a specified column, or the type of a specified column.

6. Which of the following is a following statement is a prepared statements?

- a) `Insert into department values(?, ?, ?)`
- b) `Insert into department values(x, x, x)`
- c) `SQLSetConnectOption(conn, SQL AUTOCOMMIT, 0)`
- d) `SQLTransact(conn, SQL ROLLBACK)`

[View Answer](#)

Answer: a

Explanation: ? is used as a placeholder whose value can be provided later.

7. Which of the following is used as the embedded SQL in COBOL?

- a) `EXEC SQL <embedded SQL statement >;`
- b) `EXEC SQL <embedded SQL statement > END-EXEC`
- c) `EXEC SQL <embedded SQL statement >`
- d) `EXEC SQL <embedded SQL statement > END EXEC;`

[View Answer](#)

Answer: b

Explanation: `EXEC SQL <embedded SQL statement >;` is normally in C.

8. Which of the following is used to distinguish the variables in SQL from the host language variables?

- a) .
- b) -
- c) :
- d) ,

[View Answer](#)

Answer: b

Explanation:

```
EXEC SQL
DECLARE c cursor FOR
SELECT ID, name
FROM student
WHERE tot_cred > :credit_amount;
```

- d) `EXEC SQL update END-SQL`

[View Answer](#)

Answer: c

Explanation: The SQL can be terminated by ; to terminate the sentence.

10. Which of the following is used to access large objects from a database ?

- a) `setBlob()`
- b) `getBlob()`
- c) `getBlob()`
- d) all of the mentioned

[View Answer](#)

Answer: d

Explanation: None.

Database Questions and Answers – Functions and Procedures

Find the error in the the above statement.

- a) Return type missing
- b) Dept_name is mismatched
- c) Reference relation is not mentioned
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Return integer should be given after create function for this particular function.

[View Answer](#)

Answer: b

Explanation: The count of the dept_name must be checked for the displaying from instructor relation.

3. Which of the following is used to input the entry and give the result in a variable in a procedure?

- a) Put and get
- b) Get and put
- c) Out and In
- d) In and out

[View Answer](#)

Answer: d

Explanation: Create procedure dept count proc(in dept name varchar(20), out d count integer). Here in and out refers to input and result of procedure.

[View Answer](#)

Answer: b

Explanation: Here the ‘Physics’ is in variable and d_count is out variable.

5. The format for compound statement is

- a) Begin end
- b) Begin atomic..... end
- c) Begin repeat
- d) Both Begin end and Begin atomic..... end

[View Answer](#)

Answer: d

Explanation: A compound statement is of the form begin . . . end, and it may contain multiple SQL statements between the begin and the end. A compound statement of the form begin atomic . . . end ensures that all the statements contained within it are executed as a single transaction.

Fill in the correct option :

- a) While Condition
- b) Until variable
- c) Until boolean expression
- d) Until 0

[View Answer](#)

Answer: c

Explanation: None.

[View Answer](#)

Answer: a

Explanation: The conditional statements supported by SQL include if-then-else statements by using this syntax. elif and elsif are not allowed.

8. A stored procedure in SQL is a _____

- a) Block of functions
- b) Group of Transact-SQL statements compiled into a single execution plan.
- c) Group of distinct SQL statements.
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: If it an atomic statement then the statements are in single transaction.

9. Temporary stored procedures are stored in _____ database.

- a) Master
- b) Model
- c) User specific
- d) Tempdb

[View Answer](#)

Answer: d

Explanation: None.

The above statements are used for

- a) Calling procedures
- b) Handling Exception
- c) Handling procedures
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: The SQL procedural language also supports the signaling of exception conditions, and declaring of handlers that can handle the exception, as in this code.

Database Questions and Answers – Triggers

1. A _____ is a special kind of a store procedure that executes in response to certain action on the table like insertion, deletion or updation of data.

- a) Procedures
- b) Triggers
- c) Functions
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Triggers are automatically generated when a particular operation takes place.

2. Triggers are supported in

- a) Delete
- b) Update
- c) Views
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: The triggers run after an insert, update or delete on a table. They are not supported for views.

3. The CREATE TRIGGER statement is used to create the trigger. THE _____ clause specifies the table name on which the trigger is to be attached. The _____ specifies that this is an AFTER INSERT trigger.

- a) for insert, on
- b) On, for insert
- c) For, insert
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: The triggers run after an insert, update or delete on a table. They are not supported for views.

4. What are the after triggers?

- a) Triggers generated after a particular operation
- b) These triggers run after an insert, update or delete on a table
- c) These triggers run after an insert, views, update or delete on a table
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: AFTER TRIGGERS can be classified further into three types as: AFTER INSERT Trigger, AFTER UPDATE Trigger, AFTER DELETE Trigger.

5. The variables in the triggers are declared using

- a) –
- b) @
- c) /
- d) /@

[View Answer](#)

Answer: b

Explanation: Example : declare @empid int; where empid is the variable.

6. The default extension for an Oracle SQL*Plus file is:

- a) .txt
- b) .pls
- c) .ora
- d) .sql

[View Answer](#)

Answer: d

Explanation: Example :None.

7. Which of the following is NOT an Oracle-supported trigger?

- a) BEFORE
- b) DURING
- c) AFTER
- d) INSTEAD OF

[View Answer](#)

Answer: b

Explanation: Example: During trigger is not possible in any database.

8. What are the different in triggers?

- a) Define, Create
- b) Drop, Comment
- c) Insert, Update, Delete
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Triggers are not possible for create, drop.

9. Triggers _____ enabled or disabled

- a) Can be
- b) Cannot be
- c) Ought to be
- d) Always

[View Answer](#)

Answer: a

Explanation: Triggers can be manipulated.

10. Which prefixes are available to Oracle triggers?

- a) : new only
- b) : old only
- c) Both :new and :old
- d) Neither :new nor :old

[View Answer](#)

Answer: c

Explanation: None.

Database Questions and Answers – Recursive Queries and Aggregation Features

1. Any recursive view must be defined as the union of two subqueries: a _____ query that is nonrecursive and a _____ query.

- a) Base, recursive
- b) Recursive, Base
- c) Base, Redundant
- d) View, Base

[View Answer](#)

Answer: a

Explanation: First compute the base query and add all the resultant tuples to the recursively defined view relation.

2. Ranking of queries is done by which of the following?

- a) Group by
- b) Order by
- c) Having
- d) Both Group by and Order by

[View Answer](#)

Answer: b

Explanation: Order by clause arranges the values in ascending or descending order where a default is ascending order.

3. In rank() function if one value is shared by two tuples then

- a) The rank order continues as counting numbers
- b) The rank order continues by leaving one rank in the middle
- c) The user specifies the order
- d) The order does not change

[View Answer](#)

Answer: b

Explanation: Example. rank() over (order by (GPA) desc).

4. The _____ function that does not create gaps in the ordering.

- a) Intense_rank()
- b) Continue_rank()
- c) Default_rank()
- d) Dense_rank()

[View Answer](#)

Answer: d

Explanation: For dense_rank() the tuples with the second highest value all get rank 2, and tuples with the third highest value get rank 3, and so on.

Inorder to give only 10 rank on the whole we should use

- a) Limit 10
- b) Upto 10
- c) Only 10
- d) Max 10

[View Answer](#)

Answer: a

Explanation: However, the limit clause does not support partitioning, so we cannot get the top n within each partition without performing ranking; further, if more than one student gets the same GPA, it is possible that one is included in the top 10, while another is excluded.

6. If there are n tuples in the partition and the rank of the tuple is r, then its _____ is defined as $(r-1)/(n-1)$.

- a) Ntile()
- b) Cum_rank
- c) Percent_rank
- d) rank()

[View Answer](#)

Answer: c

Explanation: Percent rank of a tuple gives the rank of the tuple as a fraction.

7. Inorder to simplify the null value confusion in the rank function we can specify

- a) Not Null
- b) Nulls last
- c) Nulls first
- d) Either Nulls last or first

[View Answer](#)

Answer: d

Explanation: select ID, rank () over (order by GPA desc nulls last) as s rank from student grades;

- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Suppose that instead of going back a fixed number of tuples, we want the window to consist of all prior years we use rows unbounded preceding.

9. The functions which construct histograms and use buckets for ranking is

- a) Rank()
- b) Newtile()
- c) Ntile()
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: For each tuple, ntile(n) then gives the number of the bucket in which it is placed, with bucket numbers starting with 1.

10. The command _____ such tables are available only within the transaction executing the query and are dropped when the transaction finishes.

- a) Create table
- b) Create temporary table
- c) Create view
- d) Create label view

[View Answer](#)

Answer: b

Explanation: None.

Database Questions and Answers – OLAP

1. OLAP stands for

- a) Online analytical processing
- b) Online analysis processing
- c) Online transaction processing
- d) Online aggregate processing

[View Answer](#)

Answer: a

Explanation: OLAP is the manipulation of information to support decision making.

2. Data that can be modeled as dimension attributes and measure attributes are called _____ data.

- a) Multidimensional
- b) Singledimensional
- c) Measured
- d) Dimensional

[View Answer](#)

Answer: a

Explanation: Given a relation used for data analysis, we can identify some of its attributes as measure attributes, since they measure some value, and can be aggregated upon. Dimension attribute define the dimensions on which measure attributes, and summaries of measure attributes, are viewed.

3. The generalization of cross-tab which is represented visually is _____ which is also called as data cube.

- a) Two dimensional cube
- b) Multidimensional cube
- c) N-dimensional cube
- d) Cuboid

[View Answer](#)

Answer: a

Explanation: Each cell in the cube is identified for the values for the three dimensional attributes.

4. The process of viewing the cross-tab (Single dimensional) with a fixed value of one attribute is

- a) Slicing
- b) Dicing
- c) Pivoting
- d) Both Slicing and Dicing

[View Answer](#)

Answer: a

Explanation: The slice operation selects one particular dimension from a given cube and provides a new sub-cube. Dice selects two or more dimensions from a given cube and provides a new sub-cube.

5. The operation of moving from finer-granularity data to a coarser granularity (by means of aggregation) is called a _____

- a) Rollup
- b) Drill down
- c) Dicing
- d) Pivoting

[View Answer](#)

Answer: a

Explanation: The opposite operation—that of moving fromcoarser-granularity data to finer-granularity data—is called a drill down.

6. In SQL the cross-tabs are created using

- a) Slice
- b) Dice
- c) Pivot
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Pivot (sum(quantity) for color in ('dark', 'pastel', 'white')).

This can be achieved by using which of the following ?

- a) group by rollup
- b) group by cubic
- c) group by
- d) none of the mentioned

[View Answer](#)

Answer: d

Explanation: 'Group by cube' is used .

8. What do data warehouses support?

- a) OLAP
- b) OLTP
- c) OLAP and OLTP
- d) Operational databases

[View Answer](#)

Answer: a

Explanation: None.

How many grouping is possible in this rollup?

- a) 8
- b) 4
- c) 2
- d) 1

[View Answer](#)

Answer: b

Explanation: { (item name, color, clothes size), (item name, color), (item name), () }.

10. Which one of the following is the right syntax for DECODE?

- a) DECODE (search, expression, result [, search, result]... [, default])
- b) DECODE (expression, result [, search, result]... [, default], search)
- c) DECODE (search, result [, search, result]... [, default], expression)
- d) DECODE (expression, search, result [, search, result]... [, default])

[View Answer](#)

Answer: d

Explanation: None.

Database Questions and Answers – Relational Algebra

1. Relational Algebra is a _____ query language that takes two relations as input and produces another relation as an output of the query.

- a) Relational
- b) Structural
- c) Procedural
- d) Fundamental

[View Answer](#)

Answer: c

Explanation: This language has fundamental and other operations which are used on relations.

2. Which of the following is a fundamental operation in relational algebra?

- a) Set intersection
- b) Natural join
- c) Assignment
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: The fundamental operations are select, project, union, set difference, Cartesian product, and rename.

3. Which of the following is used to denote the selection operation in relational algebra?

- a) Pi (Greek)
- b) Sigma (Greek)
- c) Lambda (Greek)
- d) Omega (Greek)

[View Answer](#)

Answer: b

Explanation: The select operation selects tuples that satisfy a given predicate.

4. For select operation the _____ appear in the subscript and the _____ argument appears in the parenthesis after the sigma.

- a) Predicates, relation
- b) Relation, Predicates
- c) Operation, Predicates
- d) Relation, Operation

[View Answer](#)

Answer: a

Explanation: None.

5. The _____ operation, denoted by $-$, allows us to find tuples that are in one relation but are not in another.

- a) Union
- b) Set-difference
- c) Différence
- d) Intersection

[View Answer](#)

Answer: b

Explanation: The expression $r - s$ produces a relation containing those tuples in r but not in s .

6. Which is a unary operation:

- a) Selection operation
- b) Primitive operation

- c) Projection operation
- d) Generalized selection

[View Answer](#)

Answer: d

Explanation: Generalization Selection takes only one argument for operation.

7. Which is a join condition contains an equality operator:

- a) Equijoins
- b) Cartesian
- c) Natural
- d) Left

[View Answer](#)

Answer: a

Explanation: None.

8. In precedence of set operators, the expression is evaluated from

- a) Left to left
- b) Left to right
- c) Right to left
- d) From user specification

[View Answer](#)

Answer: b

Explanation: The expression is evaluated from left to right according to the precedence.

9. Which of the following is not outer join?

- a) Left outer join
- b) Right outer join
- c) Full outer join
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: The FULL OUTER JOIN keyword combines the result of both LEFT and RIGHT joins.

10. The assignment operator is denoted by

- a) \rightarrow
- b) \leftarrow
- c) =
- d) ==

[View Answer](#)

Answer: b

Explanation: The result of the expression to the right of the \leftarrow is assigned to the relation variable on the left of the \leftarrow .

Database Questions and Answers – Tuple Relational Calculus and Domain Relational Calculus

3.

```
{t | ∃ s ∈ instructor (t[name] = s[name]
  ∧ ∃ u ∈ department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

1. Find the ID, name, dept name, salary for instructors whose salary is greater than \$80,000 .

- a) {t | t ∈ instructor ∧ t[salary] > 80000}
- b) ∃ t ∈ r (Q(t))
- c) {t | ∃ s ∈ instructor (t[ID] = s[ID] ∧ s[salary] > 80000)}
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: This expression is in tuple relational format.

3.

```
{t | ∃ s ∈ instructor (t[name] = s[name]
  ∧ ∃ u ∈ department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

2. A query in the tuple relational calculus is expressed as:

- a) {t | P0 | t}
- b) {P(t) | t}
- c) {t | P(t)}
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: The tuple relational calculus, is a nonprocedural query language. It describes the desired information without giving a specific procedure for obtaining that information.

3.

```
{t | ∃ s ∈ instructor (t[name] = s[name]
  ∧ ∃ u ∈ department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

Which of the following best describes the query?

- a) Finds the names of all instructors whose department is in the Watson building
- b) Finds the names of all department is in the Watson building
- c) Finds the name of the department whose instructor and building is Watson
- d) Returns the building name of all the departments

[View Answer](#)

Answer: a

Explanation: This query has two “there exists” clauses in our tuple-relational-calculus expression, connected by and (\wedge).

3.

```
{t | ∃ s ∈ instructor (t[name] = s[name]
  ∧ ∃ u ∈ department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

4. Which of the following symbol is used in the place of except?

- a) ^

b) V

c) \neg

d) \sim

[View Answer](#)

Answer: c

Explanation: The query $\neg P$ negates the value of P.

3.

```
{t | ∃ s ε instructor (t[name] = s[name]
  ∧ ∃ u ε department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

5. "Find all students who have taken all courses offered in the Biology department." The expressions that matches this sentence is :

- a) $\exists t \in r (Q(t))$
- b) $\forall t \in r (Q(t))$
- c) $\neg t \in r (Q(t))$
- d) $\sim t \in r (Q(t))$

[View Answer](#)

Answer: b

Explanation: \forall is used denote "for all" in SQL.

3.

```
{t | ∃ s ε instructor (t[name] = s[name]
  ∧ ∃ u ε department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

6. Which of the following is the comparison operator in tuple relational calculus

- a) \Rightarrow
- b) $=$
- c) ϵ
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: The comparison operators are ($<$, \leq , $=$, \geq , $>$).

3.

```
{t | ∃ s ε instructor (t[name] = s[name]
  ∧ ∃ u ε department (u[dept name] = s[dept name]
  ∧ u[building] = "Watson"))}
```

7. An expression in the domain relational calculus is of the form

- a) $\{P(x_1, x_2, \dots, x_n) | <x_1, x_2, \dots, x_n>\}$
- b) $\{x_1, x_2, \dots, x_n | <x_1, x_2, \dots, x_n>\}$
- c) $\{x_1, x_2, \dots, x_n | x_1, x_2, \dots, x_n\}$
- d) $\{<x_1, x_2, \dots, x_n> | P(x_1, x_2, \dots, x_n)\}$

[View Answer](#)

Answer: d

Explanation: Here x_1, x_2, \dots, x_n represent domain variables. P represents a formula composed of atoms, as was the case in the tuple relational calculus.

3.

```
{t | ∃ s ε instructor (t[name] = s[name]
```

$\Lambda \exists u \in \text{department} (u[\text{dept name}] = s[\text{dept name}] \wedge u[\text{building}] = \text{"Watson"}) \}$

[View Answer](#)

Answer: b

Explanation: None.

3.

$\{t \mid \exists s \in \text{instructor} (t[\text{name}] = s[\text{name}] \wedge \exists u \in \text{department} (u[\text{dept name}] = s[\text{dept name}] \wedge u[\text{building}] = \text{"Watson"})) \}$

9. In domain relational calculus “there exist” can be expressed as

- a) $(P1(x))$
- b) $(P1(x)) \exists x$
- c) $\forall x (P1(x))$
- d) $\exists x (P1(x))$

[View Answer](#)

Answer: d

Explanation: \exists is used to denote “some” values in relational calculus.

3.

$\{t \mid \exists s \in \text{instructor} (t[\text{name}] = s[\text{name}] \wedge \exists u \in \text{department} (u[\text{dept name}] = s[\text{dept name}] \wedge u[\text{building}] = \text{"Watson"})) \}$

10. A set of possible data values is called

- a) Attribute
- b) Degree
- c) Tuple
- d) Domain

[View Answer](#)

Answer: d

Explanation: None.

Database Questions and Answers – The Entity-Relationship Model

1. An _____ is a set of entities of the same type that share the same properties, or attributes.

- a) Entity set
- b) Attribute set
- c) Relation set
- d) Entity model

[View Answer](#)

Answer: a

Explanation: An entity is a “thing” or “object” in the real world that is distinguishable from all other objects.

2. Entity is a _____

- a) Object of relation
- b) Present working model
- c) Thing in real world
- d) Model of relation

[View Answer](#)

Answer: c

Explanation: For example, each person in a university is an entity.

3. The descriptive property possessed by each entity set is _____

- a) Entity
- b) Attribute
- c) Relation
- d) Model

[View Answer](#)

Answer: b

Explanation: Possible attributes of the instructor entity set are ID, name, dept name, and salary.

4. The function that an entity plays in a relationship is called that entity’s _____

- a) Participation
- b) Position
- c) Role
- d) Instance

[View Answer](#)

Answer: c

Explanation: A relationship is an association among several entities.

5. The attribute *name* could be structured as an attribute consisting of first name, middle initial, and last name. This type of attribute is called

- a) Simple attribute
- b) Composite attribute
- c) Multivalued attribute
- d) Derived attribute

[View Answer](#)

Answer: b

Explanation: Composite attributes can be divided into subparts (that is, other attributes).

6. The attribute AGE is calculated from DATE_OF_BIRTH. The attribute AGE is

- a) Single valued
- b) Multi valued

- c) Composite
- d) Derived

[View Answer](#)

Answer: d

Explanation: The value for this type of attribute can be derived from the values of other related attributes or entities.

7. Not applicable condition can be represented in relation entry as

- a) NA
- b) 0
- c) NULL
- d) Blank Space

[View Answer](#)

Answer: c

Explanation: NULL always represents that the value is not present.

8. Which of the following can be a multivalued attribute?

- a) Phone_number
- b) Name
- c) Date_of_birth
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Name and Date_of_birth cannot hold more than 1 value.

9. Which of the following is a single valued attribute

- a) Register_number
- b) Address
- c) SUBJECT_TAKEN
- d) Reference

[View Answer](#)

Answer: a

Explanation: None.

10. In a relation between the entities the type and condition of the relation should be specified. That is called as _____ attribute.

- a) Descriptive
- b) Derived
- c) Recursive
- d) Relative

[View Answer](#)

Answer: a

Explanation: Consider the entity sets student and section, which participate in a relationship set takes. We may wish to store a descriptive attribute grade with the relationship to record the grade that a student got in the class.

Database Questions and Answers – Constraints

1. _____ express the number of entities to which another entity can be associated via a relationship set.

- a) Mapping Cardinality
- b) Relational Cardinality
- c) Participation Constraints
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Mapping cardinality is also called as cardinality ratio.

2. An entity in A is associated with at most one entity in B, and an entity in B is associated with at most one entity in A. This is called as

- a) One-to-many
- b) One-to-one
- c) Many-to-many
- d) Many-to-one

[View Answer](#)

Answer: b

Explanation: Here one entity in one set is related to one entity in other set.

3. An entity in A is associated with at most one entity in B. An entity in B, however, can be associated with any number (zero or more) of entities in A.

- a) One-to-many
- b) One-to-one
- c) Many-to-many
- d) Many-to-one

[View Answer](#)

Answer: d

Explanation: Here more than one entity in one set is related to one entity in other set.

4. Data integrity constraints are used to:

- a) Control who is allowed access to the data
- b) Ensure that duplicate records are not entered into the table
- c) Improve the quality of data entered for a specific property
- d) Prevent users from changing the values stored in the table

[View Answer](#)

Answer: c

Explanation: The data entered will be in a particular cell (i.e., table column).

5. Establishing limits on allowable property values, and specifying a set of acceptable, predefined options that can be assigned to a property are examples of

- a) Attributes
- b) Data integrity constraints
- c) Method constraints
- d) Referential integrity constraints

[View Answer](#)

Answer: b

Explanation: Only particular value satisfying the constraints are entered in the column.

6. Which of the following can be addressed by enforcing a referential integrity constraint?

- a) All phone numbers must include the area code
- b) Certain fields are required (such as the email address, or phone number) before the record is accepted
- c) Information on the customer must be known before anything can be sold to that customer
- d) When entering an order quantity, the user must input a number and not some text (i.e., 12 rather than 'a dozen')

[View Answer](#)

Answer: c

Explanation: None.

7. _____ is a special type of integrity constraint that relates two relations & maintains consistency across the relations.

- a) Entity Integrity Constraints
- b) Referential Integrity Constraints
- c) Domain Integrity Constraints
- d) Domain Constraints

[View Answer](#)

Answer: b

Explanation: None.

8. Which one of the following uniquely identifies the elements in the relation?

- a) Secondary Key
- b) Primary key
- c) Foreign key
- d) Composite key

[View Answer](#)

Answer: b

Explanation: Primary key checks for not null and uniqueness constraint.

9. Drop Table cannot be used to drop a table referenced by a _____ constraint.

- a) Local Key
- b) Primary Key
- c) Composite Key
- d) Foreign Key

[View Answer](#)

Answer: d

Explanation: Foreign key is used when primary key of one relation is used in another relation.

10. _____ is preferred method for enforcing data integrity

- a) Constraints
- b) Stored Procedure
- c) Triggers
- d) Cursors

[View Answer](#)

Answer: a

Explanation: Constraints are specified to restrict entries in the relation.

Database Questions and Answers – Entity-Relationship Diagrams

1. Which of the following gives a logical structure of the database graphically?

- a) Entity-relationship diagram
- b) Entity diagram
- c) Database diagram
- d) Architectural representation

[View Answer](#)

Answer: a

Explanation: E-R diagrams are simple and clear—qualities that may well account in large part for the widespread use of the E-R model.

2. The entity relationship set is represented in E-R diagram as

- a) Double diamonds
- b) Undivided rectangles
- c) Dashed lines
- d) Diamond

[View Answer](#)

Answer: d

Explanation: Dashed lines link attributes of a relationship set to the relationship set.

3. The Rectangles divided into two parts represents

- a) Entity set
- b) Relationship set
- c) Attributes of a relationship set
- d) Primary key

[View Answer](#)

Answer: a

Explanation: The first part of the rectangle, contains the name of the entity set. The second part contains the names of all the attributes of the entity set.

4. Consider a directed line(->) from the relationship set advisor to both entity sets instructor and student. This indicates _____ cardinality

- a) One to many
- b) One to one
- c) Many to many
- d) Many to one

[View Answer](#)

Answer: b

Explanation: This indicates that an instructor may advise at most one student, and a student may have at most one advisor.

5. We indicate roles in E-R diagrams by labeling the lines that connect _____ to _____

- a) Diamond , diamond
- b) Rectangle, diamond
- c) Rectangle, rectangle
- d) Diamond, rectangle

[View Answer](#)

Answer: d

Explanation: Diamond represents a relationship set and rectangle represents a entity set.

6. An entity set that does not have sufficient attributes to form a primary key is termed a _____

- a) Strong entity set

- b) Variant set
- c) Weak entity set
- d) Variable set

[View Answer](#)

Answer: c

Explanation: An entity set that has a primary key is termed a strong entity set.

7. For a weak entity set to be meaningful, it must be associated with another entity set, called the

- a) Identifying set
- b) Owner set
- c) Neighbour set
- d) Strong entity set

[View Answer](#)

Answer: a

Explanation: Every weak entity must be associated with an identifying entity; that is, the weak entity set is said to be existence dependent on the identifying entity set. The identifying entity set is said to own the weak entity set that it identifies. It is also called as owner entity set.

8. Weak entity set is represented as

- a) Underline
- b) Double line
- c) Double diamond
- d) Double rectangle

[View Answer](#)

Answer: c

Explanation: An entity set that has a primary key is termed a strong entity set.

9. If you were collecting and storing information about your music collection, an album would be considered a(n) _____

- a) Relation
- b) Entity
- c) Instance
- d) Attribute

[View Answer](#)

Answer: b

Explanation: An entity set is a logical container for instances of an entity type and instances of any type derived from that entity type.

10. What term is used to refer to a specific record in your music database; for instance; information stored about a specific album?

- a) Relation
- b) Instance
- c) Table
- d) Column

[View Answer](#)

Answer: b

Explanation: The environment of database is said to be an instance. A database instance or an ‘instance’ is made up of the background processes needed by the database.

Database Questions and Answers – Reduction to Relational Schemas

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

1. Which one of the following can be treated as a primary key in teaches relation?

- a) Id
- b) Semester
- c) Sec_id
- d) Year

[View Answer](#)

Answer: a

Explanation: Here Id is the only attribute that has to have a unique entry.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

2. The primary key in the section relation is

- a) Course_id
- b) Sec_id
- c) Both Course_id and Sec_id
- d) All the attributes

[View Answer](#)

Answer: c

Explanation: Both the entries has unique entry.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

Which of the following Id is selected for the following query?

- a) 1003
- b) 1001
- c) None
- d) Error message appears

[View Answer](#)

Answer: d

Explanation: The value CS-101 matches the Course_id but not Id.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

Which of the following Id are displayed?

- a) 1003
- b) 1001
- c) Both 1003 and 1001
- d) Error message appears

[View Answer](#)

Answer: c

Explanation: Two rows are select in the above query.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

5. The query which selects the Course_id ‘CS-101’ from the section relation is

- a) Select Course_id from section where Building = ‘Richard’;
- b) Select Course_id from section where Year = ‘2009’;
- c) Select Course_id from teaches where Building = ‘Packyard’;
- d) Select Course_id from section where Sec_id = ‘3’;

[View Answer](#)

Answer: b

Explanation: The year ‘2009’ should be selected from the section relation.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

Which of the following has an error in the above create table for the relation section

- a) Primary key (course id, sec id, semester, year)
- b) Foreign key (course id) references course
- c) Year numeric (4,0)
- d) Building numeric (15)

[View Answer](#)

Answer: d

Explanation: It should be replaced by Year Building varchar (15).

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

7. The relation with primary key can be created using

- a) Create table instructor (Id, Name)
- b) Create table instructor (Id, Name, primary key(name))
- c) Create table instructor (Id, Name, primary key (Id))
- d) Create table instructor (Id unique, Name)

[View Answer](#)

Answer: c

Explanation: The value Name cannot be a primary key.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

8. How can the values in the relation teaches be deleted?

- a) Drop table teaches;
- b) Delete from teaches;
- c) Purge table teaches;
- d) Delete from teaches where Id =’Null’;

[View Answer](#)

Answer: b

Explanation: Delete table cleans the entry from the table.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

9. In the above teaches relation " Select * from teaches where Year = '2010'" displays how many rows?

- a) 2
- b) 4
- c) 5
- d) 1

[View Answer](#)

Answer: a

Explanation: There are two tuples with the year is 2009.

3.

```
SELECT * FROM teaches WHERE Sec_id = 'CS-101';
```

10. The relation changes can be got back using _____ command

- a) Flashback
- b) Purge
- c) Delete
- d) Getback

[View Answer](#)

Answer: a

Explanation: Purge deletes the table and delete cleans the table entry.

ayzom.com

Database Questions and Answers – Entity-Relationship Design Issues

1. Let us consider *phone_number* ,which can take single or several values . Treating *phone_number* as an _____ permits instructors to have several phone numbers (including zero) associated with them.

- a) Entity
- b) Attribute
- c) Relation
- d) Value

[View Answer](#)

Answer: a

Explanation: Treating a phone as an attribute *phone_number* implies that instructors have precisely one phone number each.

2. The total participation by entities is represented in E-R diagram as

- a) Dashed line
- b) Double line
- c) Double rectangle
- d) Circle

[View Answer](#)

Answer: b

Explanation: It is used to represent the relation between several attributes.

3. Given the basic ER and relational models, which of the following is INCORRECT?

- a) An attribute of an entity can have more than one value
- b) An attribute of an entity can be composite
- c) In a row of a relational table, an attribute can have more than one value
- d) In a row of a relational table, an attribute can have exactly one value or a NULL value

[View Answer](#)

Answer: c

Explanation: It is possible to have several values for a single attribute provide it is a multi-valued attribute.

4. Which of the following indicates the maximum number of entities that can be involved in a relationship?

- a) Minimum cardinality
- b) Maximum cardinality
- c) ERD
- d) Greater Entity Count

[View Answer](#)

Answer: b

Explanation: In SQL (Structured Query Language), the term cardinality refers to the uniqueness of data values contained in a particular column (attribute) of a database table.

5. In E-R diagram generalization is represented by

- a) Ellipse
- b) Dashed ellipse
- c) Rectangle
- d) Triangle

[View Answer](#)

Answer: d

Explanation: Ellipse represents attributes, rectangle represents entity.

6. What is a relationship called when it is maintained between two entities?

- a) Unary
- b) Binary
- c) Ternary
- d) Quaternary

[View Answer](#)

Answer: b

Explanation: Binary word usually represents two attributes.

7. Which of the following is a low level operator?

- a) Insert
- b) Update
- c) Delete
- d) Directory

[View Answer](#)

Answer: d

Explanation: Directory is a low level word on in file system.

8. Key to represent relationship between tables is called

- a) Primary key
- b) Secondary Key
- c) Foreign Key
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Primary key of one relation used as an attribute in another relation is called foreign key.

9. A window into a portion of a database is

- a) Schema
- b) View
- c) Query
- d) Data dictionary

[View Answer](#)

Answer: b

Explanation: View is a logical portion of a database which is needed by some users.

10. A primary key is combined with a foreign key creates

- a) Parent-Child relationship between the tables that connect them
- b) Many to many relationship between the tables that connect them
- c) Network model between the tables that connect them
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Using the two relationships mother and father provides us a record of a child's mother, even if we are not aware of the father's identity; a null value would be required if the ternary relationship parent is used. Using binary relationship sets is preferable in this case.

Database Questions and Answers – Extended E-R Features

1. The entity set person is classified as student and employee. This process is called _____

- a) Generalization
- b) Specialization
- c) Inheritance
- d) Constraint generalization

[View Answer](#)

Answer: b

Explanation: The process of designating subgroupings within an entity set is called specialization.

2. Which relationship is used to represent a specialization entity?

- a) ISA
- b) AIS
- c) ONIS
- d) WHOIS

[View Answer](#)

Answer: a

Explanation: In terms of an E-R diagram, specialization is depicted by a hollow arrow-head pointing from the specialized entity to the other entity.

3. The refinement from an initial entity set into successive levels of entity subgroupings represents a _____ design process in which distinctions are made explicit.

- a) Hierarchy
- b) Bottom-up
- c) Top-down
- d) Radical

[View Answer](#)

Answer: c

Explanation: The design process may also proceed in a bottom-up manner, in which multiple entity sets are synthesized into a higher-level entity set on the basis of common features.

4. There are similarities between the instructor entity set and the secretary entity set in the sense that they have several attributes that are conceptually the same across the two entity sets: namely, the identifier, name, and salary attributes. This process is called

- a) Commonality
- b) Specialization
- c) Generalization
- d) Similarity

[View Answer](#)

Answer: c

Explanation: Generalization is used to emphasize the similarities among lower-level entity sets and to hide the differences.

5. If an entity set is a lower-level entity set in more than one ISA relationship, then the entity set has

- a) Hierarchy
- b) Multilevel inheritance
- c) Single inheritance
- d) Multiple inheritance

[View Answer](#)

Answer: d

Explanation: The attributes of the higher-level entity sets are said to be inherited by the lower-level entity sets.

6. A _____ constraint requires that an entity belong to no more than one lower-level entity set.

- a) Disjointness
- b) Uniqueness
- c) Special
- d) Relational

[View Answer](#)

Answer: a

Explanation: For example, student entity can satisfy only one condition for the student type attribute; an entity can be either a graduate student or an undergraduate student, but cannot be both.

7. Consider the employee work-team example, and assume that certain employees participate in more than one work team. A given employee may therefore appear in more than one of the team entity sets that are lower level entity sets of employee. Thus, the generalization is

- a) Overlapping
- b) Disjointness
- c) Uniqueness
- d) Relational

[View Answer](#)

Answer: a

Explanation: In overlapping generalizations, the same entity may belong to more than one lower-level entity set within a single generalization.

8. The completeness constraint may be one of the following: Total generalization or specialization, Partial generalization or specialization. Which is the default?

- a) Total
- b) Partial
- c) Should be specified
- d) Cannot be determined

[View Answer](#)

Answer: b

Explanation: Partial generalization or specialization – Some higher-level entities may not belong to any lower-level entity set.

9. Functional dependencies are a generalization of

- a) Key dependencies
- b) Relation dependencies
- c) Database dependencies
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: The subclasses are combined to form the superclass.

10. Which of the following is another name for a weak entity?

- a) Child
- b) Owner
- c) Dominant
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: A parent may be called as a strong entity.

Database Questions and Answers – Querying database part-1 DDL

1. Which is the main relation which is used in the university database which is referenced by all other relation of the university?

- a) Teaches
- b) Course
- c) Department
- d) Section

[View Answer](#)

Answer: c

Explanation: Department is the only relation which forms the main part of the university database.

2. The department relation has the an entry budget whose type has to be replaced by

- a) Varchar (20)
- b) Varchar2 (20)
- c) Numeric (12,2)
- d) Numeric

[View Answer](#)

Answer: c

Explanation: Department is the only relation which forms the main part of the university database.

3. In the course relation, the title field should throw an error in case of any missing title. The command to be added in title is

- a) Unique
- b) Not null
- c) 0
- d) Null

[View Answer](#)

Answer: b

Explanation: By specifying not null the value cannot be left blank.

4. In the above DDL command the foreign key entries are got by using the keyword

- a) References
- b) Key reference
- c) Relating
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: References (table_name) give the prior table name for the entry.

5. Identify the error in the section relation

- a) No error
- b) Year numeric (4,0)
- c) Building varchar (15)
- d) Sec_id varchar (8)

[View Answer](#)

Answer: a

Explanation: The building and the sec_id have varchar values and year is of numeric type. So no such errors are found in the relation.

Identify the output of the query given

- a) Row(s) inserted
- b) Error in ID of insert

- c) Error in Name of insert
- d) Error in Salary of the insert

[View Answer](#)

Answer: b

Explanation: The varchar(5) value cannot hold the entry 100202.

7. Which of the following can be used as a primary key entry of the instructor relation.

- a) DEPT_NAME
- b) NAME
- c) ID
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: The value ID can only be primary key unlike dept_name which is used as a foreign key.

8. In the section relation which of the following is used as a foreign key?

- a) Course_id
- b) Course_id,sec_id
- c) Room_number
- d) Course_id,sec_id,room_number

[View Answer](#)

Answer: a

Explanation: Course_id is the only field which is present in the course relation.

9. In order to include an attribute Name to the teaches relation which of the following command is used?

- a) Alter table teaches include Name;
- b) Alter table teaches add Name;
- c) Alter table teaches add Name varchar;
- d) Alter table teaches add Name varchar(20);

[View Answer](#)

Answer: d

Explanation: The form of the alter table command is

alter table r add A D;

where r is the name of an existing relation, A is the name of the attribute to be added, and D is the type of the added attribute.

10. To replace the relation section with some other relation the initial step to be carried out is

- a) Delete section;
- b) Drop section;
- c) Delete from section;
- d) Replace section new_table ;

[View Answer](#)

Answer: b

Explanation: Dropping the table drops all the references to that table.

Database Questions and Answers – Querying database part-2 DML

1. Which of the following command is used to display the departments of the instructor relation?

- a) Select * from instructor where Dept_name = Finance;
- b) Select * from instructor ;
- c) Select dept_name from instructor;
- d) Select dept_name for instructor where Name=Jackson;

[View Answer](#)

Answer: c

Explanation: Only one field is necessary for the query and where clause is not needed for the selection.

2. How can we select the elements which have common Dept_name in both the relation ?

- a) Select * from instructor i , course c where i.Dept_name=c.Dept_name;
- b) Select Dept name from instructor ,Course ;
- c) Select * from instructor i , course c ;
- d) Select Dept_name from instructor where Dept_name = NULL;

[View Answer](#)

Answer: a

Explanation: Here only the common elements are displayed .

3. Select distinct Dept_name from instructor ;

How many row(s) are displayed ?

- a) 4
- b) 3
- c) 5
- d) Error

[View Answer](#)

Answer: a

Explanation: Distinct keyword eliminates the the common Dept_name .

d) Insert into course values(12111,Emma,NeuroScience,200000);

[View Answer](#)

Answer: c

Explanation: The values have to be inserted into both the relations to be intact .

5. If a person all the people in Music department gets fired which of the following has to be performed on the instructor relation ?

- a) Delete Dept_name=Music in instructor;
- b) Delete from instructor where Dept_name=Music;
- c) Remove Dept_name= Music
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: Delete from *table_name* where *condition* .

What will be displayed as the value of name for the above query?

- a) Hayley
- b) Jackson
- c) Hayley and Crick
- d) Crick

[View Answer](#)

Answer: d

Explanation: Only the greatest salary in Comp.Sci dept is selected for the query.

How many rows are selected ?

- a) 3
- b) 4
- c) 2
- d) 1

[View Answer](#)

Answer: d

Explanation: This displays the names of instructors with salary greater than that of some (at least one) instructor in the Biology department .

[View Answer](#)

Answer: a

Explanation: % is used to indicate that some characters may appear .

9. Which function is used to find the count of distinct departments?

- a) Dist
- b) Distinct
- c) Count
- d) Count,Dist

[View Answer](#)

Answer: a

Explanation: Count (distinct ID) is the correct usage.

10. Which function is used to identify the title with Least scope?

- a) Min(Credits)
- b) Max(Credits)
- c) Min(title)
- d) Min(Salary)

[View Answer](#)

Answer: a

Explanation: Max is used to find the highest element and Min is used to find the lowest element.

Database Questions and Answers – Atomic Domains

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

1. A domain is _____ if elements of the domain are considered to be indivisible units.

- a) Atomic
- b) Subatomic
- c) Substructure
- d) Subset

[View Answer](#)

Answer: a

Explanation: A set of names is an example of a nonatomic value.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

2. Identify the composite attributes

- a) Salary
- b) Credits
- c) Section_id
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Composite attributes, such as an attribute address with component attributes street, city, state, and zip have nonatomic domains.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

- a) i, ii and iii are in 3NF and iv is in BCNF
- b) i and ii are in BCNF and iii and iv are in 3NF
- c) All are in 3NF
- d) All are in BCNF

[View Answer](#)

Answer: d

Explanation: One of the more desirable normal forms that we can obtain is Boyce–Codd normal form (BCNF). It eliminates all redundancy that can be discovered based on functional dependencies.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

4. Which one is based on multi-valued dependency:

- a) First
- b) Second
- c) Third
- d) Fourth

[View Answer](#)

Answer: d

Explanation: One of the more desirable normal forms that we can obtain is Boyce–Codd normal form (BCNF). It eliminates all redundancy that can be discovered based on functional dependencies.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

5. If a relation is in BCNF, then it is also in

- a) 1 NF
- b) 2 NF
- c) 3 NF
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: Third normal form (3NF) relaxes this constraint slightly by allowing certain nontrivial functional dependencies whose left side is not a superkey.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

6. If every non-key attribute is functionally dependent primary key, then the relation will be in

- a) First normal form
- b) Second normal form
- c) Third form
- d) Fourth normal form

[View Answer](#)

Answer: b

Explanation: Third normal form (3NF) relaxes this constraint slightly by allowing certain nontrivial functional dependencies whose left side is not a superkey.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. $R(A, B)$ WITH productions { $A \rightarrow B$ }
- ii. $R(A, B)$ WITH productions { $B \rightarrow A$ }
- iii. $R(A, B)$ WITH productions { $A \rightarrow B, B \rightarrow A$ }
- iv. $R(A, B, C)$ WITH productions { $A \rightarrow B, B \rightarrow A, AB \rightarrow C$ }

7. If an attribute of a composite key is dependent on an attribute of the other composite key, a normalization called _____ is needed.

- a) DKNF
- b) BCNF
- c) Fourth
- d) Third

[View Answer](#)

Answer: b

Explanation: BCNF eliminates all redundancy that can be discovered based on functional dependencies.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. R(A, B) WITH productions { A --> B }
- ii. R(A, B) WITH productions { B --> A }
- iii. R(A, B) WITH productions {A -> B, B --> A }
- iv. R(A, B, C) WITH productions {A -->B, B --> A, AB --> C }

8. The term for information that describes what type of data is available in a database is:

- a) Data dictionary
- b) data repository
- c) Index data
- d) Metadata

[View Answer](#)

Answer: d

Explanation: Meta data is generally data about a data.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. R(A, B) WITH productions { A --> B }
- ii. R(A, B) WITH productions { B --> A }
- iii. R(A, B) WITH productions {A -> B, B --> A }
- iv. R(A, B, C) WITH productions {A -->B, B --> A, AB --> C }

9. A data type that creates unique numbers for key columns in Microsoft Access is:

- a) Autonumber
- b) Boolean
- c) Sequential key
- d) Sequential number

[View Answer](#)

Answer: a

Explanation: This can be taken as a primary key.

3. Consider the relation given below and ind the maximum normal form applicable to them

- i. R(A, B) WITH productions { A --> B }
- ii. R(A, B) WITH productions { B --> A }
- iii. R(A, B) WITH productions {A -> B, B --> A }
- iv. R(A, B, C) WITH productions {A -->B, B --> A, AB --> C }

10. A dependency exist between two columns when

- a) Together they constitute a composite key for the table
- b) Knowing the value in one column determines the value stored in another column
- c) The table is in 3NF
- d) Together they constitute a foreign key

[View Answer](#)

Answer: a

Explanation: Given a set F of functional dependencies on a schema, we can prove that certain other functional dependencies also hold on the schema.

Database Questions & Answers – Normal Forms

1. In the _____ normal form, a composite attribute is converted to individual attributes.

- a) First
- b) Second
- c) Third
- d) Fourth

[View Answer](#)

Answer: a

Explanation: The first normal form is used to eliminate the duplicate information.

2. A table on the many side of a one to many or many to many relationship must:

- a) Be in Second Normal Form (2NF)
- b) Be in Third Normal Form (3NF)
- c) Have a single attribute key
- d) Have a composite key

[View Answer](#)

Answer: d

Explanation: The relation in second normal form is also in first normal form and no partial dependencies on any column in primary key.

3. Tables in second normal form (2NF):

- a) Eliminate all hidden dependencies
- b) Eliminate the possibility of a insertion anomalies
- c) Have a composite key
- d) Have all non key fields depend on the whole primary key

[View Answer](#)

Answer: a

Explanation: The relation in second normal form is also in first normal form and no partial dependencies on any column in primary key.

4. Which-one of the following statements about normal forms is FALSE?

- a) BCNF is stricter than 3 NF
- b) Lossless, dependency -preserving decomposition into 3 NF is always possible
- c) Loss less, dependency – preserving decomposition into BCNF is always possible
- d) Any relation with two attributes is BCNF

[View Answer](#)

Answer: c

Explanation: We say that the decomposition is a lossless decomposition if there is no loss of information by replacing r (R) with two relation schemas r1(R1) and r2(R2).

5. Functional Dependencies are the types of constraints that are based on_____

- a) Key
- b) Key revisited
- c) Superset key
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Key is the basic element needed for the constraints.

6. Which is a bottom-up approach to database design that design by examining the relationship between attributes:

- a) Functional dependency

- b) Database modeling
- c) Normalization
- d) Decomposition

[View Answer](#)

Answer: c

Explanation: Normalisation is the process of removing redundancy and unwanted data.

7. Which forms simplifies and ensures that there are minimal data aggregates and repetitive groups:

- a) 1NF
- b) 2NF
- c) 3NF
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: The first normal form is used to eliminate the duplicate information.

8. Which forms has a relation that possesses data about an individual entity:

- a) 2NF
- b) 3NF
- c) 4NF
- d) 5NF

[View Answer](#)

Answer: c

Explanation: A Table is in 4NF if and only if, for every one of its non-trivial multivalued dependencies $X \twoheadrightarrow Y$, X is a superkey—that is, X is either a candidate key or a superset thereof.

9. Which forms are based on the concept of functional dependency:

- a) 1NF
- b) 2NF
- c) 3NF
- d) 4NF

[View Answer](#)

Answer: c

Explanation: The table is in 3NF if every non-prime attribute of R is non-transitively dependent (i.e. directly dependent) on every superkey of R.

For any pincode, there is only one city and state. Also, for given street, city and state, there is just one pincode. In normalization terms, empdt1 is a relation in

- a) 1 NF only
- b) 2 NF and hence also in 1 NF
- c) 3NF and hence also in 2NF and 1NF
- d) BCNF and hence also in 3NF, 2NF and 1NF

[View Answer](#)

Answer: b

Explanation: The relation in second normal form is also in first normal form and no partial dependencies on any column in primary key.

Database Questions and Answers – Functional-Dependency Theory

3. The relation employee(ID, name, street, Credit, street, city, salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

1. We can use the following three rules to find logically implied functional dependencies. This collection of rules is called

- a) Axioms
- b) Armstrong's axioms
- c) Armstrong
- d) Closure

[View Answer](#)

Answer: b

Explanation: By applying these rules repeatedly, we can find all of F^+ , given F .

3. The relation employee(ID, name, street, Credit, street, city, salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

2. Which of the following is not Armstrong's Axiom?

- a) Reflexivity rule
- b) Transitivity rule
- c) Pseudotransitivity rule
- d) Augmentation rule

[View Answer](#)

Answer: c

Explanation: It is possible to use Armstrong's axioms to prove that Pseudotransitivity rule is sound.

3. The relation employee(ID, name, street, Credit, street, city, salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

This type of decomposition is called

- a) Lossless decomposition
- b) Lossless-join decomposition
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Lossy-join decomposition is the decomposition used here .

3. The relation employee(ID, name, street, Credit, street, city, salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

This comes under

- a) Lossy-join decomposition
- b) Lossy decomposition
- c) Lossless-join decomposition
- d) Both Lossy and Lossy-join decomposition

[View Answer](#)

Answer: d

Explanation: Lossy-join decomposition is the decomposition used here .

3. The relation employee(ID,name,street,Credit,street,city,salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

5. There are two functional dependencies with the same set of attributes on the left side of the arrow:

A->BC

A->B

This can be combined as

- a) A->BC
- b) A->B
- c) B->C
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: This can be computed as the canonical cover.

3. The relation employee(ID,name,street,Credit,street,city,salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

The number of superkeys of R is:

- a) 2
- b) 7
- c) 10
- d) 12

[View Answer](#)

Answer: c

Explanation: A superkey is a combination of columns that uniquely identifies any row within a relational database management system (RDBMS) table.

3. The relation employee(ID,name,street,Credit,street,city,salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

Which, if any, of the two queries above will correctly (in SQL2) get the desired set of employee ID's?

- a) Both I and II
- b) I only
- c) II only
- d) Neither I nor II

[View Answer](#)

Answer: a

Explanation: The query can be satisfied by any of the two options.

3. The relation employee(ID,name,street,Credit,street,city,salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

- a) 2
- b) 4
- c) 6

d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: The SQL NATURAL JOIN is a type of EQUI JOIN and is structured in such a way that, columns with same name of associate tables will appear once only.

3. The relation employee(ID,name,street,Credit,street,city,salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

Then which of the following is the most restrictive, correct condition on the value of m?

- a) $m = \min(r,s)$
- b) $0 \leq m \leq r + s$
- c) $\min(r,s) \leq m \leq \max(r,s)$
- d) $0 \leq m \leq \min(r,s)$

[View Answer](#)

Answer: d

Explanation: The value of m must lie between the min value of r and s and 0.

3. The relation employee(ID,name,street,Credit,street,city,salary) is decomposed into

```
employee1 (ID, name)
employee2 (name, street, city, salary)
```

Which of the following is not a key?

- a) A
- b) E
- c) B, C
- d) D

[View Answer](#)

Answer: c

Explanation: Here the keys are not formed by B and C.

Database Questions and Answers – Algorithms for Decomposition

1. A relation is in _____ if an attribute of a composite key is dependent on an attribute of other composite key.

- a) 2NF
- b) 3NF
- c) BCNF
- d) 1NF

[View Answer](#)

Answer: b

Explanation: A relation is in 3 NF if an attribute of a composite key is dependent on an attribute of other composite key. (If an attribute of a composite key is dependent on an attribute of other composite key then the relation is not in BCNF, hence it has to be decomposed.).

2. What are the desirable properties of a decomposition

- a) Partition constraint
- b) Dependency preservation
- c) Redundancy
- d) Security

[View Answer](#)

Answer: b

Explanation: Lossless join and dependency preserving are the two goals of the decomposition.

3. R (A,B,C,D) is a relation. Which of the following does not have a lossless join dependency preserving BCNF decomposition?

- a) A->B, B->CD
- b) A->B, B->C, C->D
- c) AB->C, C->AD
- d) A->BCD

[View Answer](#)

Answer: d

Explanation: This relation gives a relation without any loss in the values.

Consider the above conditions. Which of the following relation holds?

- a) Course id-> title, dept name, credits
- b) Title-> dept name, credits
- c) Dept name-> credits
- d) Cannot be determined

[View Answer](#)

Answer: a

Explanation: Here course id is not a superkey. Thus, class is not in BCNF.

5. The algorithm that takes a set of dependencies and adds one schema at a time, instead of decomposing the initial schema repeatedly is

- a) BCNF algorithm
- b) 2NF algorithm
- c) 3NF synthesis algorithm
- d) 1NF algorithm

[View Answer](#)

Answer: c

Explanation: The result is not uniquely defined, since a set of functional dependencies can have more than one canonical cover, and, further, in some cases, the result of the algorithm depends on the order in which it considers the dependencies in Fc.

6. The functional dependency can be tested easily on the materialized view, using the constraints _____.

- a) Primary key
- b) Null
- c) Unique
- d) Both Null and Unique

[View Answer](#)

Answer: d

Explanation: Primary key contains both unique and not null constraints.

7. Which normal form is considered adequate for normal relational database design?

- a) 2NF
- b) 5NF
- c) 4NF
- d) 3NF

[View Answer](#)

Answer: d

Explanation: A relational database table is often described as “normalized” if it is in the Third Normal Form because most of the 3NF tables are free of insertion, update, and deletion anomalies.

8. Relation R with an associated set of functional dependencies, F, is decomposed into BCNF. The redundancy (arising out of functional dependencies) in the resulting set of relations is

- a) Zero
- b) More than zero but less than that of an equivalent 3NF decomposition
- c) Proportional to the size of F^+
- d) Indeterminate

[View Answer](#)

Answer: b

Explanation: Redundancy in BCNF is low when compared to 3NF. For more details on BCNF.

in terms of normalization, this table is in

- a) 1NF
- b) 2NF
- c) 3NF
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Since the primary key is not given we have to derive the primary key of the table. Using the closure set of attributes we get the primary key as $(F1, F2)$. From functional dependencies, “ $F1 \rightarrow F3, F2 \rightarrow F4$ ”, we can see that there is partial functional dependency therefore it is not in 1NF. Hence the table is in 1NF.

The relation schema R is

- a) in BCNF
- b) in 3NF, but not in BCNF
- c) in 2NF, but not in 3NF
- d) not in 2NF

[View Answer](#)

Answer: d

Explanation: From the closure set of attributes we can see that the key for the relation is AB. The FD $B \rightarrow G$ is a partial dependency, hence it is not in 2NF.

Database Questions and Answers – Multivalued Dependencies

1. The normal form which satisfies multivalued dependencies and which is in BCNF is

- a) 4 NF
- b) 3 NF
- c) 2 NF
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Fourth normal form is more restrictive than BCNF.

2. Which of the following is a tuple-generating dependencies?

- a) Functional dependency
- b) Equality-generating dependencies
- c) Multivalued dependencies
- d) Non-functional dependency

[View Answer](#)

Answer: c

Explanation: Multivalued dependencies, do not rule out the existence of certain tuples. Instead, they require that other tuples of a certain form be present in the relation.

3. The main task carried out in the _____ is to remove repeating attributes to separate tables.

- a) First Normal Form
- b) Second Normal Form
- c) Third Normal Form
- d) Fourth Normal Form

[View Answer](#)

Answer: a

Explanation: Multivalued dependencies, do not rule out the existence of certain tuples. Instead, they require that other tuples of a certain form be present in the relation.

4. Which of the normal form is based on multivalued dependencies?

- a) First
- b) Second
- c) Third
- d) Fourth

[View Answer](#)

Answer: d

Explanation: Multivalued dependencies, do not rule out the existence of certain tuples. Instead, they require that other tuples of a certain form be present in the relation.

5. Which forms has a relation that possesses data about an individual entity?

- a) 2NF
- b) 3NF
- c) 4NF
- d) 5NF

[View Answer](#)

Answer: c

Explanation: A Table is in 4NF if and only if, for every one of its non-trivial multivalued dependencies $X \twoheadrightarrow Y$, X is a superkey —that is, X is either a candidate key or a superset thereof.

6. If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from one of the following sources.

- a) A many-to-many relationship set
- b) A multivalued attribute of an entity set
- c) A one-to-many relationship set
- d) Both A many-to-many relationship set and A multivalued attribute of an entity set

[View Answer](#)

Answer: d

Explanation: For a many-to-many relationship set each related entity set has its own schema and there is an additional schema for the relationship set. For a multivalued attribute, a separate schema is created consisting of that attribute and the primary key of the entity set.

7. Which of the following has each related entity set has its own schema and there is an additional schema for the relationship set?

- a) A many-to-many relationship set
- b) A multivalued attribute of an entity set
- c) A one-to-many relationship set
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from this source.

8. In which of the following, a separate schema is created consisting of that attribute and the primary key of the entity set.

- a) A many-to-many relationship set
- b) A multivalued attribute of an entity set
- c) A one-to-many relationship set
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from this source.

9. Fifth Normal form is concerned with

- a) Functional dependency
- b) Multivalued dependency
- c) Join dependency
- d) Domain-key

[View Answer](#)

Answer: c

Explanation: If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from this source.

10. In 2NF

- a) No functional dependencies (FDs) exist
- b) No multivalued dependencies (MVDs) exist
- c) No partial FDs exist
- d) No partial MVDs exist

[View Answer](#)

Answer: c

Explanation: If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from this source.

Database Questions and Answers – Database Design Process

1. _____ can help us detect poor E-R design.

- a) Database Design Process
- b) E-R Design Process
- c) Relational scheme
- d) Functional dependencies

[View Answer](#)

Answer: d

Explanation: For eg., Suppose an instructor entity set had attributes dept name and dept address, and there is a functional dependency dept name \rightarrow dept address.

2. If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from one of the following sources.

- a) A many-to-many relationship set
- b) A multivalued attribute of an entity set
- c) A one-to-many relationship set
- d) Both A many-to-many relationship set and A multivalued attribute of an entity set

[View Answer](#)

Answer: d

Explanation: For a many-to-many relationship set each related entity set has its own schema and there is an additional schema for the relationship set. For a multivalued attribute, a separate schema is created consisting of that attribute and the primary key of the entity set.

3. Which of the following has each related entity set has its own schema and there is an additional schema for the relationship set.

- a) A many-to-many relationship set
- b) A multivalued attribute of an entity set
- c) A one-to-many relationship set
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from this source.

4. In which of the following, a separate schema is created consisting of that attribute and the primary key of the entity set.

- a) A many-to-many relationship set
- b) A multivalued attribute of an entity set
- c) A one-to-many relationship set
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: If a multivalued dependency holds and is not implied by the corresponding functional dependency, it usually arises from this source.

5. Suppose the user finds the usage of *room number* and *phone number* in a relational schema there is confusion. This is reduced by

- a) Unique-role assumption
- b) Unique-key assignment
- c) Role intergral assignment
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: A desirable feature of a database design is the unique-role assumption, which means that each attribute name has a unique meaning in the database.

6. What is the best way to represent the attributes in a large database?

- a) Relational-and
- b) Concatenation
- c) Dot representation
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: Example inst sec and student sec.

7. Designers use which of the following to tune the performance of systems to support time-critical operations?

- a) Denormalization
- b) Redundant optimization
- c) Optimization
- d) Realization

[View Answer](#)

Answer: a

Explanation: The process of taking a normalized schema and making it nonnormalized is called denormalization.

8. In the schema (dept name, size) we have relations *total inst 2007*, *total inst 2008*. Which dependency have lead to this relation ?

- a) Dept name, year->size
- b) Year->size
- c) Dept name->size
- d) Size->year

[View Answer](#)

Answer: a

Explanation: The process of taking a normalized schema and making it nonnormalized is called denormalization.

9. Relation *dept year(dept name, total inst 2007, total inst 2008, total inst 2009)*. Here the only functional dependencies are from dept name to the other attributes. This relation is in

- a) Fourth NF
- b) BCNF
- c) Third NF
- d) Second NF

[View Answer](#)

Answer: b

Explanation: BCNF has only one normal form.

10. Thus a _____ of course data gives the values of all attributes, such as title and department, of all courses at a particular point in time.

- a) Instance
- b) Snapshot
- c) Both Instance and Snapshot
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: We use the term snapshot of data to mean the value of the data at a particular point in time.

11. Representations such as the in the dept year relation, with one column for each value of an attribute, are called _____ they are widely used in spreadsheets and reports and in data analysis tools.

- a) Cross-tabs
- b) Snapshot
- c) Both Cross-tabs and Snapshot
- d) All of the mentioned

Akhilesh Yadav | [Linkedin.com/in/arki7n](https://www.linkedin.com/in/arki7n) | [instagram.com/arki7n](https://www.instagram.com/arki7n)

[View Answer](#)

Answer: a

Explanation: SQL includes features to convert data from a normal relational representation to a crosstab.

ayzom.com

Database Questions and Answers – Application Programs and User Interfaces

1. An approach to website design with the emphasis on converting visitors to outcomes required by the owner is referred to as:

- a) Web usability
- b) Persuasion
- c) Web accessibility
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: In computing, graphical user interface is a type of user interface that allows users to interact with electronic devices.

2. A method of modelling and describing user tasks for an interactive application is referred to as:

- a) Customer journey
- b) Primary persona
- c) Use case
- d) Web design persona

[View Answer](#)

Answer: c

Explanation: The actions in GUI are usually performed through direct manipulation of the graphical elements.

3. Information architecture influences:

- a) Answer choice
- b) Site structure
- c) Labeling
- d) Navigation design

[View Answer](#)

Answer: b

Explanation: The actions in GUI are usually performed through direct manipulation of the graphical elements.

4. Also known as schematics, a way of illustrating the layout of an individual webpage is a:

- a) Wireframe
- b) Sitemap
- c) Card sorting
- d) Blueprint

[View Answer](#)

Answer: a

Explanation: An application programming interface specifies how some software components should interact with each other.

5. A graphical or text depiction of the relationship between different groups of content on a website is referred to as a:

- a) Wireframe
- b) Blueprint
- c) Sitemap
- d) Card sorting

[View Answer](#)

Answer: c

Explanation: An application programming interface specifies how some software components should interact with each other.

6. Blueprints are intended to:

- a) Prototype of the screen layout showing navigation and main design elements
- b) Show the grouping of pages and user journeys

- c) Indicate the structure of a site during site design and as a user feature
- d) Prototype typical customer journeys or clickstreams through a website

[View Answer](#)

Answer: c

Explanation: A blueprint is a reproduction of a technical drawing, documenting an architecture or an engineering design, using a contact print process.

7. Storyboards are intended to:

- a) Indicate the structure of a site during site design and as a user feature
- b) Prototype of the screen layout showing navigation and main design elements
- c) Integrate consistently available components on the webpage (e.g. navigation, search boxes)
- d) Prototype typical customer journeys or click streams through a website

[View Answer](#)

Answer: d

Explanation: An application programming interface specifies how some software components should interact with each other.

8. Which of the following occupies boot record of hard and floppy disks and activated during computer startup?

- a) Worm
- b) Boot sector virus
- c) Macro virus
- d) Virus

[View Answer](#)

Answer: b

Explanation: A blueprint is a reproduction of a technical drawing, documenting an architecture or an engineering design, using a contact print process.

9. A graphical or text depiction of the relationship between different groups of content on a website is a:

- a) Page template
- b) Wireframe
- c) Site map
- d) Cascading style sheet (CSS)

[View Answer](#)

Answer: c

Explanation: In computing, graphical user interface is a type of user interface that allows users to interact with electronic devices.

10. Which of the following is a description of information organization schemes?

- a) Minimising the number of clicks needed to access relevant content
- b) Providing an overall design to a site consistent with the positioning of the products and services
- c) The menu options chosen to group and categorize information
- d) Providing specific content and services appropriate to different audience members

[View Answer](#)

Answer: c

Explanation: In computing, graphical user interface is a type of user interface that allows users to interact with electronic devices.

Database Questions and Answers – Web Fundamentals

1. Which of the following is a valid uniform resource locator?

- a) <http://www.acm.org/sigmod>
- b) www.google.com
- c) www.ann.in
- d) <http://www.acm.org/sigmod/>

[View Answer](#)

Answer: a

Explanation: A uniform resource locator (URL) is a globally unique name for each document that can be accessed on the Web.

2. <http://www.google.com/search?q=silberschatz>

In the above URL which one is the argument which is used for processing of the URL?

- a) google
- b) google.com
- c) search
- d) q=silberschatz

[View Answer](#)

Answer: d

Explanation: Argument is always placed after ? symbol.

3. HTTP defines two ways in which values entered by a user at the browser can be sent to the Web server. The _____ method encodes the values as part of the URL.

- a) Post
- b) Get
- c) Read
- d) Argument

[View Answer](#)

Answer: b

Explanation: For example, if the Google search page used a form with an input parameter named q with the get method, and the user typed in the string “silberschatz” and submitted the form, the browser would request the following URL from the Web server: <http://www.google.com/search?q=silberschatz>.

4. A _____ is a program running on the server machine, which accepts requests from a Web browser and sends back results in the form of HTML documents.

- a) HTML
- b) HTTP
- c) Web Server
- d) Web browser

[View Answer](#)

Answer: c

Explanation: The browser and Web server communicate via HTTP. Web servers provide powerful features, beyond the simple transfer of documents.

5. The application program typically communicates with a database server, through _____ or other protocols, in order to get or store data.

- a) JDBC
- b) ODBC
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: The common gateway interface (CGI) standard defines how the Web server communicates with application programs.

6. This extra information is usually maintained in the form of a _____ at the client.

- a) Cookie
- b) History
- c) Remainder
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: A cookie is simply a small piece of text containing identifying information and with an associated name.

7. Which of the following is not true about HTML ?

- a) <meta>...</meta>
- b) <meta.../>
- c) <metadata>...</metadata>
- d) <metadata name="" />

[View Answer](#)

Answer: b

Explanation: Meta data is the data about data which is included in the meta data tag.

8. Html code contains:

- a) Tags
- b) Attributes
- c) Elements
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: <> are tags, size is a attribute.

9. Html document must always be saved with:

- a) .html
- b) .htm
- c) .doc
- d) Both .html & .htm

[View Answer](#)

Answer: d

Explanation: .doc is used only for the word document format.

10. How many levels of headings are in html:

- a) 2
- b) 7
- c) 6
- d) 4

[View Answer](#)

Answer: c

Explanation: The heading levels are h1,h2,h3,h4,h5,h6.

Database Questions and Answers – Servlets and JSP

1. The Java _____ specification defines an application programming interface for communication between the Web server and the application program.

- a) Servlet
- b) Server
- c) Program
- d) Randomize

[View Answer](#)

Answer: a

Explanation: Servlets are commonly used to generate dynamic responses to HTTP requests.

2. The doGet() method in the example extracts values of the parameter's type and number by using _____

- a) request.getParameter()
- b) request.setParameter()
- c) responce.getParameter()
- d) responce.getAttribute()

[View Answer](#)

Answer: a

Explanation: These methods uses these values to run a query against a database.

3. How many JDBC driver types does Sun define?

- a) One
- b) Two
- c) Three
- d) Four

[View Answer](#)

Answer: d

Explanation: JBDB.DriverManager.getConnection() is used to get the connection to the database.

4. Which JDBC driver Type(s) can be used in either applet or servlet code?

- a) Both Type 1 and Type 2
- b) Both Type 1 and Type 3
- c) Both Type 3 and Type 4
- d) Type 4 only

[View Answer](#)

Answer: c

Explanation: In a Type 3 driver, a three-tier approach is used to accessing databases. The JDBC clients use standard network sockets to communicate with an middleware application server. In a Type 4 driver, a pure Java-based driver that communicates directly with vendor's database through socket connection.

5. What MySQL property is used to create a surrogate key in MySQL?

- a) UNIQUE
- b) SEQUENCE
- c) AUTO_INCREMENT
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: A surrogate key in a database is a unique identifier for either an entity in the modeled world or an object in the database.

6. A JSP is transformed into a(n):

- a) Java applet
- b) Java servlet
- c) Either 1 or 2 above
- d) Neither 1 nor 2 above

[View Answer](#)

Answer: b

Explanation: Servlets are commonly used to generate dynamic responses to HTTP requests.

7. Which JDBC driver Type(s) is(are) the JDBC-ODBC bridge?

- a) Type 1
- b) Type 2
- c) Type 3
- d) Type 4

[View Answer](#)

Answer: a

Explanation: In a Type 1 driver, a JDBC bridge is used to access ODBC drivers installed on each client machine.

8. What programming language(s) or scripting language(s) does Java Server Pages (JSP) support?

- a) VBScript only
- b) Jscript only
- c) Java only
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: JSP primarily uses Java for certain codes.

9. What is bytecode?

- a) Machine-specific code
- b) Java code
- c) Machine-independent code
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Java bytecode is the form of instructions that the Java virtual machine executes. Each bytecode opcode is one byte in length, although some require parameters, resulting in some multi-byte instructions.

10. Where is metadata stored in MySQL?

- a) In the MySQL database metadata
- b) In the MySQL database metasql
- c) In the MySQL database mysql
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Metadata contains data about other data which is given in the <meta>...</meta> tags.

Database Questions and Answers – Application Architectures

1. Which of the following is true for Seeheim model?
a) Presentation is abstracted from dialogue and Application
b) Presentation and Dialogue is abstracted from Application
c) Presentation and Application is abstracted from Dialogue
d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Presentation is abstracted from dialogue and application.

2. Which of the unit operation is used in Model view controller?

- a) Is a Decomposition
- b) Part Whole Decomposition
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Part whole decomposition is applied to MVC.

3. Memory address refers to the successive memory words and the machine is called as _____

- a) word addressable
- b) byte addressable
- c) bit addressable
- d) Terra byte addressable

[View Answer](#)

Answer: a

Explanation: Part whole decomposition is applied to MVC.

4. Which layer deals which deals with user interaction is called _____ layer.

- a) Business logic
- b) Presentation
- c) User interaction
- d) Data access

[View Answer](#)

Answer: b

Explanation: The single application may have several different versions of this layer, corresponding to distinct kinds of interfaces such as Web browsers, and user interfaces of mobile phones, which have much smaller screens.

5. The _____ layer, which provides a high-level view of data and actions on data.

- a) Business logic
- b) Presentation
- c) User interaction
- d) Data access

[View Answer](#)

Answer: a

Explanation: The single application may have several different versions of this layer, corresponding to distinct kinds of interfaces such as Web browsers, and user interfaces of mobile phones, which have much smaller screens.

6. The _____ layer, which provides the interface between the business-logic layer and the underlying database.

- a) Business logic
- b) Presentation
- c) User interaction
- d) Data access

[View Answer](#)

Answer: d

Explanation: Many applications use an object-oriented language to code the business-logic layer, and use an object-oriented model of data, while the underlying database is a relational database.

7. The _____ system is widely used for mapping from Java objects to relations.

- a) Hibernate
- b) Object oriented
- c) Objective
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: In Hibernate, the mapping from each Java class to one or more relations is specified in a mapping file.

8. Which among the following are the functions that any system with a user interface must provide?

- a) Presentation
- b) Dialogue
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Presentation and Application are the functions that any system with a user interface must provide.

9. Which of the following is the main task accomplished by the user?

- a) Compose a document
- b) Create a spread sheet
- c) Send mail
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: All of the mentioned are the main task accomplished by the user.

10. What are the portability concerns founded in Seeheim model?

- a) Replacing the presentation toolkit
- b) Replacing the application toolkit
- c) Replacing the dialogue toolkit
- d) Replacing the presentation & application toolkit

[View Answer](#)

Answer: d

Explanation: The portability concerns founded in Seeheim model are- Replacing the presentation toolkit and Replacing the application toolkit.

Database Questions and Answers – Rapid Application Development

1. Which schema object instructs Oracle to connect to a remotely access an object of a database?

- a) Sequence
- b) Remote link
- c) Database link
- d) Data link

[View Answer](#)

Answer: d

Explanation: A database link (DBlink) is a definition of how to establish a connection from one Oracle database to another.

2. DML changes are

- a) Insert
- b) Update
- c) Create
- d) Both Insert and Update

[View Answer](#)

Answer: d

Explanation: Create is a DDL operation.

3. Which of the following object types below cannot be replicated?

- a) Data
- b) Trigger
- c) View
- d) Sequence

[View Answer](#)

Answer: d

Explanation: Sequence is a series of items which is like a unique index.

4. How to force a log switch?

- a) By using ALTER SYSTEM LOG
- b) By using ALTER SYSTEM SWITCH LOGFILE
- c) By using ALTER SYSTEM SWITCH LOGS
- d) By using ALTER SYS LOGFILES

[View Answer](#)

Answer: b

Explanation: ALTER SYSTEM ARCHIVE LOG CURRENT is the best practice for production backup scripts with RMAN..

- a) 100 / 0.15
- b) quantity – 100
- c) 35*20
- d) 0.15-35

[View Answer](#)

Answer: a

Explanation: According to the precedence of expression as in BODMAS the expression evaluated.

6. The ORDER BY clause can only be used in

- a) SELECT queries
- b) INSERT queries
- c) GROUP BY queries

d) HAVING queries

[View Answer](#)

Answer: a

Explanation: `SELECT column_name,column_name
FROM table_name
ORDER BY column_name,column_name ASC|DESC;`

7. Which of the following rule below are categories of an index?

- a) Column and Functional
- b) Multiple Column and functional
- c) Column, Multiple Column and functional
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: The CREATE INDEX statement is used to create indexes in tables.

8. What is the purpose of SMON background process?

- a) Performs crash recovery when a failed instance starts up again
- b) Performs recovery when a user process fails
- c) Writes redo log entries to disk
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: SMON (System MONitor) is an Oracle background process created when you start a database instance.

9. Which of the following queries are legal?

- a) `SELECT deptno, count(deptno) FROM emp GROUP BY ename;`
- b) `SELECT deptno, count(deptno), job FROM emp GROUP BY deptno;`
- c) `SELECT deptno, avg(sal) FROM emp;`
- d) `SELECT deptno, avg(sal) FROM emp GROUP BY deptno;`

[View Answer](#)

Answer: d

Explanation: For aggregate functions group by clause is necessary.

10. Which of the following queries displays the sum of all employee salaries for those employees not making commission, for each job, including only those sums greater than 2500?

- a) `select job, sum(sal) from emp where sum(sal) > 2500 and comm is null;`
- b) `select job, sum(sal) from emp where comm is null group by job having sum(sal) > 2500;`
- c) `select job, sum(sal) from emp where sum(sal) > 2500 and comm is null group by job;`
- d) `select job, sum(sal) from emp group by job having sum(sal) > 2500 and comm is not null;`

[View Answer](#)

Answer: b

Explanation: For aggregate functions group by clause is necessary.

Database Questions and Answers – Application Performance

1. The indirect change of the values of a variable in one module by another module is called
 - a) Internal change
 - b) Inter-module change
 - c) Side effect
 - d) Side-module update

[View Answer](#)

Answer: c

Explanation: The module of the search tree and the flow is directed by its values.

2. Which of the following data structure is not linear data structure?
 - a) Arrays
 - b) Linked lists
 - c) Arrays & Linked lists
 - d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Both array and linked lists are in data structure concepts.

3. Which of the following data structure is linear data structure?
 - a) Trees
 - b) Graphs
 - c) Arrays
 - d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Tree and graphs are not linear.

4. Which of the following criterion is NOT written using the proper syntax?
 - a) "Haris"
 - b) <500
 - c) NO VALUE
 - d) Between #1/1/2000# and #12/31/2000#

[View Answer](#)

Answer: c

Explanation: NO VALUE cannot be specified.

5. The operation of processing each element in the list is known as
 - a) Sorting
 - b) Merging
 - c) Inserting
 - d) Traversal

[View Answer](#)

Answer: d

Explanation: There are several types of traversals.

6. Finding the location of the element with a given value is:
 - a) Traversal
 - b) Search

- c) Sort
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Search is performed by traversing through the tree.

7. Arrays are best data structures

- a) For relatively permanent collections of data
- b) For the size of the structure and the data in the structure are constantly changing
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: The operator tree has a tree like format where the evaluation starts from root of the tree.

8. Linked lists are best suited

- a) For relatively permanent collections of data
- b) For the size of the structure and the data in the structure are constantly changing
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: A linked list is a data structure consisting of a group of nodes which together represent a sequence.

9. Each array declaration need not give, implicitly or explicitly, the information about

- a) The name of array
- b) The data type of array
- c) The first data from the set to be stored
- d) The index set of the array

[View Answer](#)

Answer: c

Explanation: The operator tree has a tree like format where the evaluation starts from root of the tree.

10. The elements of an array are stored successively in memory cells because

- a) By this way computer can keep track only the address of the first element and the addresses of other elements can be calculated
- b) The architecture of computer memory does not allow arrays to store other than serially
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Memory is always allotted in order.

Database Questions and Answers – Application Security

1. In _____ attacks, the attacker manages to get an application to execute an SQL query created by the attacker.

- a) SQL injection
- b) SQL
- c) Direct
- d) Application

[View Answer](#)

Answer: a

Explanation: Application security has to deal with several security threats and issues beyond those handled by SQL authorization.

2. A Web site that allows users to enter text, such as a comment or a name, and then stores it and later display it to other users, is potentially vulnerable to a kind of attack called a _____ attack.

- a) Two-factor authentication
- b) Cross-site request forgery
- c) Cross-site scripting
- d) Cross-site scoring scripting

[View Answer](#)

Answer: c

Explanation: In such an attack, a malicious user enters code written in a client-side scripting language such as JavaScript or Flash instead of entering a valid name or comment.

3. _____ is an attack which forces an end user to execute unwanted actions on a web application in which he/she is currently authenticated.

- a) Two-factor authentication
- b) Cross-site request forgery
- c) Cross-site scripting
- d) Cross-site scoring scripting

[View Answer](#)

Answer: b

Explanation: Cross-site request forgery, also known as a one-click attack or session riding and abbreviated as CSRF or XSRF.

4. Many applications use _____ where two independent factors are used to identify a user.

- a) Two-factor authentication
- b) Cross-site request forgery
- c) Cross-site scripting
- d) Cross-site scoring scripting

[View Answer](#)

Answer: a

Explanation: The two factors should not share a common vulnerability.

5. Even with two-factor authentication, users may still be vulnerable to _____ attacks.

- a) Radian
- b) Cross attack
- c) scripting
- d) Man-in-the-middle

[View Answer](#)

Answer: d

Explanation: In such attacks, a user attempting to connect to the application is diverted to a fake Web site, which accepts the password from the user, and uses it immediately to authenticate to the original application.

6. A single _____ further allows the user to be authenticated once, and multiple applications can then verify the user's identity through an authentication service without requiring reauthentication.

- a) OpenID
- b) Sign-on system
- c) Security Assertion Markup Language (SAML)
- d) Virtual Private Database (VPD)

[View Answer](#)

Answer: b

Explanation: Once the user logged in at one site, he does not have to enter his user name and password at other sites that use the same single sign-on service.

7. The _____ is a standard for exchanging authentication and authorization information between different security domains, to provide cross-organization single sign-on.

- a) OpenID
- b) Sign-on system
- c) Security Assertion Markup Language (SAML)
- d) Virtual Private Database (VPD)

[View Answer](#)

Answer: c

Explanation: The user's password and other authentication factors are never revealed to the application, and the user need not register explicitly with the application.

8. The _____ standard is an alternative for single sign-on across organizations, and has seen increasing acceptance in recent years.

- a) OpenID
- b) Single-site system
- c) Security Assertion Markup Language (SAML)
- d) Virtual Private Database (VPD)

[View Answer](#)

Answer: a

Explanation: The user's password and other authentication factors are never revealed to the application, and the user need not register explicitly with the application.

9. _____ allows a system administrator to associate a function with a relation; the function returns a predicate that must be added to any query that uses the relation.

- a) OpenID
- b) Single-site system
- c) Security Assertion Markup Language (SAML)
- d) Virtual Private Database (VPD)

[View Answer](#)

Answer: d

Explanation: Some database systems provide mechanisms for fine-grained authorization.

10. VPD provides authorization at the level of specific tuples, or rows, of a relation, and is therefore said to be a _____ mechanism.

- a) Row-level authorization
- b) Column-level authentication
- c) Row-type authentication
- d) Authorization security

[View Answer](#)

Answer: a

Explanation: Oracle Virtual Private Database (VPD) allows a system administrator to associate a function with a relation.

Database Questions and Answers – Encryption and Its Applications

1. _____ is widely used today for protecting data in transit in a variety of applications such as data transfer on the Internet, and on cellular phone networks.

- a) Encryption
- b) Data mining
- c) Internet Security
- d) Architectural security

[View Answer](#)

Answer: a

Explanation: Encryption is also used to carry out other tasks, such as authentication.

2. In a database where the encryption is applied the data is cannot be handled by the unauthorised user without

- a) Encryption key
- b) Decryption key
- c) Primary key
- d) Authorised key

[View Answer](#)

Answer: b

Explanation: Even if the message is intercepted by an enemy, the enemy, not knowing the key, will not be able to decrypt and understand the message.

3. Which of the following is not a property of good encryption technique?

- a) Relatively simple for authorized users to encrypt and decrypt data
- b) Decryption key is extremely difficult for an intruder to determine
- c) Encryption depends on a parameter of the algorithm called the encryption key
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Here a, b and c are the properties have to be present in a good design of an encryption technique.

4. In which of the following encryption key is used to encrypt and decrypt the data?

- a) Public key
- b) Private key
- c) Symmetric key
- d) Asymmetric key

[View Answer](#)

Answer: c

Explanation: In public-key (also known as asymmetric-key) encryption techniques, there are two different keys, the public key and the private key, used to encrypt and decrypt the data.

5. Encryption of small values, such as identifiers or names, is made complicated by the possibility of _____

- a) Dictionary attacks
- b) Database attacks
- c) Minor attacks
- d) Random attacks

[View Answer](#)

Answer: a

Explanation: This happens when particularly if the encryption key is publicly available.

6. Which one of the following uses a 128bit round key to encrypt the data using XOR and use it in reverse to decrypt it?

- a) Round key algorithm
- b) Public key algorithm
- c) Advanced Encryption Standard
- d) Asymmetric key algorithm

[View Answer](#)

Answer: c

Explanation: The standard is based on the Rijndael algorithm.

7. Which of the following requires no password travel across the internet?

- a) Readable system
- b) Manipulation system
- c) Challenge-response system
- d) Response system

[View Answer](#)

Answer: c

Explanation: The database system sends a challenge string to the user. The user encrypts the challenge string using a secret password as encryption key and then returns the result. The database system can verify the authenticity of the user by decrypting the string with the same secret password and checking the result with the original challenge string.

8. Assymmetric Encryption: Why can a message encrypted with the Public Key only be decrypted with the receiver's appropriate Private Key?

- a) Not true, the message can also be decrypted with the Public Key
- b) A so called "one way function with back door" is applied for the encryption
- c) The Public Key contains a special function which is used to encrypt the message and which can only be reversed by the appropriate Private Key
- d) The encrypted message contains the function for decryption which identifies the Private Key

[View Answer](#)

Answer: b

Explanation: An one-way function is a function which a computer can calculate quickly, but whose reversal would last months or years. An one-way function with back door can be reversed with the help of a couple of additional information (the back door), but scarcely without this information. The information for the back door is contained in the private Key.

9. Which is the largest disadvantage of symmetric Encryption?

- a) More complex and therefore more time-consuming calculations
- b) Problem of the secure transmission of the Secret Key
- c) Less secure encryption function
- d) Isn't used any more

[View Answer](#)

Answer: b

Explanation: As there is only one key in the symmetrical encryption, this must be known by both sender and recipient and this key is sufficient to decrypt the secret message. Therefore it must be exchanged between sender and receiver in such a manner that an unauthorized person can in no case take possession of it.

10. Which is the principle of the encryption using a key?

- a) The key indicates which function is used for encryption. Thereby it is more difficult to decrypt an intercepted message as the function is unknown
- b) The key contains the secret function for encryption including parameters. Only a password can activate the key
- c) All functions are public, only the key is secret. It contains the parameters used for the encryption resp. decryption
- d) The key prevents the user of having to reinstall the software at each change in technology or in the functions for encryption

[View Answer](#)

Answer: b

Explanation: The encoding of a message is calculated by an algorithm. If always the same algorithm would be used, it would be easy to crack

Akhilesh Yadav | [Linkedin.com/in/arki7n](https://www.linkedin.com/in/arki7n) | [instagram.com/arki7n](https://www.instagram.com/arki7n)

intercepted messages. However, it isn't possible to invent a new algorithm whenever the old one was cracked, therefore the possibility to parameterize algorithms is needed and this is the assignment of the key.

ayzom.com

Visit <https://ayzom.com> | <https://t.me/arki7n>

Database Questions and Answers – Physical Storage Media

1. Which of the following is a physical storage media?

- a) Tape Storage
- b) Optical Storage
- c) Flash memory
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: The storage media are classified by the speed with which data can be accessed, by the cost per unit of data to buy the medium, and by the medium's reliability.

2. The _____ is the fastest and most costly form of storage, which is relatively small; its use is managed by the computer system hardware.

- a) Cache
- b) Disk
- c) Main memory
- d) Flash memory

[View Answer](#)

Answer: a

Explanation: Cache storage is easy to access because it is closer to the processor.

3. Which of the following stores several gigabytes of data but usually lost when power failure?

- a) Flash memory
- b) Disk
- c) Main memory
- d) Secondary memory

[View Answer](#)

Answer: c

Explanation: The contents of main memory are usually lost if a power failure or system crash occurs.

4. The flash memory storage used are

- a) NOR Flash
- b) OR Flash
- c) AND Flash
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: NAND flash has a much higher storage capacity for a given cost, and is widely used for data storage in devices such as cameras, music players, and cell phones.

5. _____ is increasingly being used in server systems to improve performance by caching frequently used data, since it provides faster access than disk, with larger storage capacity than main memory.

- a) Flash memory
- b) Disk
- c) Main memory
- d) Secondary memory

[View Answer](#)

Answer: a

Explanation: Flash memory is of two types – NAND and NOR.

6. Which is the cheapest memory device in terms of costs/ bit?

- a) Semiconductor memory
- b) Magnetic disks
- c) Compact disks
- d) Magnetic tapes

[View Answer](#)

Answer: c

Explanation: Compact disk is used for easy storage at lower cost.

7. The primary medium for the long-term online storage of data is the _____ where the entire database is stored on magnetic disk.

- a) Semiconductor memory
- b) Magnetic disks
- c) Compact disks
- d) Magnetic tapes

[View Answer](#)

Answer: b

Explanation: The system must move the data from disk to main memory so that they can be accessed.

8. Optical disk _____ systems contain a few drives and numerous disks that can be loaded into one of the drives automatically (by a robot arm) on demand.

- a) Tape Storage
- b) Jukebox
- c) Flash memory
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: The most popular form of optical disks are CD and DVD.

9. There are “record-once” versions of the compact disk and digital video disk, which can be written only once; such disks are also called _____ disks.

- a) Write-once, read-many (WORM)
- b) CD-R
- c) DVD-W
- d) CD-ROM

[View Answer](#)

Answer: a

Explanation: There are also “multiple-write” versions of compact disk (called CD-RW) and digital video disk (DVD-RW, DVD+RW, and DVD-RAM), which can be written multiple times.

10. Tape storage is referred to as _____ storage.

- a) Direct-access
- b) Random-access
- c) Sequential-access
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Tape storage is used primarily for backup and archival data.

Database Questions and Answers – Magnetic Disk and Flash Storage

1. In magnetic disk _____ stores information on a sector magnetically as reversals of the direction of magnetization of the magnetic material.

- a) Read-write head
- b) Read-assemble head
- c) Head-disk assemblies
- d) Disk arm

[View Answer](#)

Answer: d

Explanation: Each side of a platter of a disk has a read-write head that moves across the platter to access different tracks.

2. A _____ is the smallest unit of information that can be read from or written to the disk.

- a) Track
- b) Spindle
- c) Sector
- d) Platter

[View Answer](#)

Answer: c

Explanation: The disk surface is logically divided into tracks, which are subdivided into sectors.

3. The disk platters mounted on a spindle and the heads mounted on a disk arm are together known as _____

- a) Read-disk assemblies
- b) Head-disk assemblies
- c) Head-write assemblies
- d) Read-read assemblies

[View Answer](#)

Answer: b

Explanation: Each side of a platter of a disk has a read-write head that moves across the platter to access different tracks.

4. The disk controller uses _____ at each sector to ensure that the data is not corrupted on data retrieval.

- a) Checksum
- b) Unit drive
- c) Read disk
- d) Readsum

[View Answer](#)

Answer: a

Explanation: A disk controller interfaces between the computer system and the actual hardware of the disk drive.

5. _____ is the time from when a read or write request is issued to when data transfer begins.

- a) Access time
- b) Average seek time
- c) Seek time
- d) Rotational latency time

[View Answer](#)

Answer: a

Explanation: To access (that is, to read or write) data on a given sector of a disk, the arm first must move so that it is positioned over the correct track, and then must wait for the sector to appear under it as the disk rotates.

6. The time for repositioning the arm is called the _____ and it increases with the distance that the arm must move.

- a) Access time

- b) Average seek time
- c) Seek time
- d) Rotational latency time

[View Answer](#)

Answer: c

Explanation: Typical seek times range from 2 to 30 milliseconds, depending on how far the track is from the initial arm position.

7. _____ is around one-half of the maximum seek time.

- a) Access time
- b) Average seek time
- c) Seek time
- d) Rotational latency time

[View Answer](#)

Answer: b

Explanation: Average seek times currently range between 4 and 10 milliseconds, depending on the disk model.

8. Once the head has reached the desired track, the time spent waiting for the sector to be accessed to appear under the head is called the

- a) Access time
- b) Average seek time
- c) Seek time
- d) Rotational latency time

[View Answer](#)

Answer: d

Explanation: Rotational speeds of disks today range from 5400 rotations per minute (90 rotations per second) up to 15,000 rotations per minute (250 rotations per second), or, equivalently, 4 milliseconds to 11.1 milliseconds per rotation.

9. In Flash memory, the erase operation can be performed on a number of pages, called an _____ at once, and takes about 1 to 2 milliseconds.

- a) Delete block
- b) Erase block
- c) Flash block
- d) Read block

[View Answer](#)

Answer: b

Explanation: The size of an erase block (often referred to as just “block” in flash literature) is usually significantly larger than the block size of the storage system.

10. Hybrid disk drives are hard-disk systems that combine magnetic storage with a smaller amount of flash memory, which is used as a cache for frequently accessed data.

- a) Hybrid drivers
- b) Disk drivers
- c) Hybrid disk drivers
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: Frequently accessed data that are rarely updated are ideal for caching in flash memory.

Database Questions and Answers – RAID

1. Which level of RAID refers to disk mirroring with block striping?

- a) RAID level 1
- b) RAID level 2
- c) RAID level 0
- d) RAID level 3

[View Answer](#)

Answer: a

Explanation: RAID level 1 refers to disk mirroring with block striping.

2. Optical disk technology uses

- a) Helical scanning
- b) DAT
- c) A laser beam
- d) RAID

[View Answer](#)

Answer: d

Explanation: Redundant Array of Inexpensive Disks.

3. With multiple disks, we can improve the transfer rate as well by _____ data across multiple disks.

- a) Striping
- b) Dividing
- c) Mirroring
- d) Dividing

[View Answer](#)

Answer: a

Explanation: Data striping consists of splitting the bits of each byte across multiple disks; such striping is called bitlevel striping.

4. Which one of the following is a Stripping technique?

- a) Byte level stripping
- b) Raid level stripping
- c) Disk level stripping
- d) Block level stripping

[View Answer](#)

Answer: d

Explanation: Block-level striping stripes blocks across multiple disks. It treats the array of disks as a single large disk, and it gives blocks logical numbers.

5. The RAID level which mirroring is done along with striping is

- a) RAID 1+0
- b) RAID 0
- c) RAID 2
- d) Both RAID 1+0 and RAID 0

[View Answer](#)

Answer: d

Explanation: Mirroring without striping can also be used with arrays of disks, to give the appearance of a single large, reliable disk.

6. Where performance and reliability are both important, RAID level _____ is used.

- a) 0

b) 1

c) 2

d) 0+1

[View Answer](#)

Answer: d

Explanation: Mirroring without striping can also be used with arrays of disks, to give the appearance of a single large, reliable disk.

7. _____ partitions data and parity among all N+1 disks, instead of storing data in N-disks and parity in one disk.

a) Block interleaved parity

b) Block interleaved distributed parity

c) Bit parity

d) Bit interleaved parity

[View Answer](#)

Answer: b

Explanation: In level 5, all disks can participate in satisfying read requests, unlike RAID level 4, where the parity disk cannot participate, so level 5 increases the total number of requests that can be met in a given amount of time.

8. Hardware RAID implementations permit _____ that is, faulty disks can be removed and replaced by new ones without turning power off.

a) Scrapping

b) Swapping

c) Hot swapping

d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Hot

swapping reduces the mean time to repair since replacement of a disk does not have to wait until a time when the system can be shut down.

9. _____ is popular for applications such as storage of log files in a database system since it offers the best write performance.

a) RAID level 1

b) RAID level 2

c) RAID level 0

d) RAID level 3

[View Answer](#)

Answer: a

Explanation: RAID level 1 refers to disk mirroring with block striping.

10. _____ which increases the number of I/O operations needed to write a single logical block, pays a significant time penalty in terms of write performance.

a) RAID level 1

b) RAID level 2

c) RAID level 5

d) RAID level 3

[View Answer](#)

Answer: a

Explanation: In level 5, all disks can participate in satisfying read requests, unlike RAID level 4, where the parity disk cannot participate, so level 5 increases the total number of requests that can be met in a given amount of time.

Database Questions and Answers – Tertiary Storage

1. Tertiary storage is built with :

- a) a lot of money
- b) unremovable media
- c) removable media
- d) secondary storage

[View Answer](#)

Answer: c

Explanation: Tertiary storage involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage device according to the system's demands; this data is often copied to secondary storage before use.

2. Operating system is responsible for

- a) disk initialization
- b) booting from disk
- c) bad-block recovery
- d) all of the mentioned

[View Answer](#)

Answer: d

Explanation: Tertiary storage involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage device according to the system's demands; this data is often copied to secondary storage before use.

3. A typical tape drive is _____ a typical disk drive.

- a) more expensive than
- b) cheaper than
- c) of the same cost as
- d) none of the mentioned

[View Answer](#)

Answer: a

Explanation: Tertiary storage involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage device according to the system's demands; this data is often copied to secondary storage before use.

4. During recovery from a failure

- a) each pair of physical block is examined
- b) specified pair of physical block is examined
- c) first pair of physical block is examined
- d) none of the mentioned

[View Answer](#)

Answer: a

Explanation: Tertiary storage involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage device according to the system's demands; this data is often copied to secondary storage before use.

5. A magneto-optic disk is :

- a) primary storage
- b) secondary storage
- c) tertiary storage
- d) none of the mentioned

[View Answer](#)

Answer: c

Explanation: Tertiary storage involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage

device according to the system's demands; this data is often copied to secondary storage before use.

6. Which of the following are the process of selecting the data storage and data access characteristics of the database?

- a) Logical database design
- b) Physical database design
- c) Testing and performance tuning
- d) Evaluation and selecting

[View Answer](#)

Answer: b

Explanation: Physical database design is the process of selecting the data storage and data access characteristics of the database.

7. The replacement of a bad block generally is not totally automatic because

- a) data in bad block cannot be replaced
- b) data in bad block is usually lost
- c) bad block does not contain any data
- d) none of the mentioned

[View Answer](#)

Answer: b

Explanation: Physical database design is the process of selecting the data storage and data access characteristics of the database.

8. Which of the following is the oldest database model?

- a) Relational
- b) Hierarchical
- c) Physical
- d) Network

[View Answer](#)

Answer: d

Explanation: Network model has data stored in a hierarchical network flow.

9. The surface area of a tape is _____ the surface area of a disk.

- a) much lesser than
- b) much larger than
- c) equal to
- d) none of the mentioned

[View Answer](#)

Answer: b

Explanation: Network model has data stored in a hierarchical network flow.

10. Which one of the following is not a secondary storage?

- a) magnetic disks
- b) magnetic tapes
- c) ram
- d) none of the mentioned

[View Answer](#)

Answer: c

Explanation: Tertiary storage involves a robotic mechanism which will mount (insert) and dismount removable mass storage media into a storage device according to the system's demands; this data is often copied to secondary storage before use.

Database Questions and Answers – File Organisations

1. Which level of RAID refers to disk mirroring with block striping?

- a) RAID level 1
- b) RAID level 2
- c) RAID level 0
- d) RAID level 3

[View Answer](#)

Answer: a

Explanation: RAID (redundant array of independent disks) is a way of storing the same data in different places (thus, redundantly) on multiple hard disks.

2. A unit of storage that can store one or more records in a hash file organization is denoted as

- a) Buckets
- b) Disk pages
- c) Blocks
- d) Nodes

[View Answer](#)

Answer: a

Explanation: A unit of storage that can store one or more records in a hash file organization is denoted as buckets.

3. The file organization which allows us to read records that would satisfy the join condition by using one block read is

- a) Heap file organization
- b) Sequential file organization
- c) Clustering file organization
- d) Hash file organization

[View Answer](#)

Answer: c

Explanation: All systems in the cluster share a common file structure via NFS, but not all disks are mounted on all other systems.

4. What are the correct features of a distributed database?

- a) Is always connected to the internet
- b) Always requires more than three machines
- c) Users see the data in one global schema.
- d) Have to specify the physical location of the data when an update is done

[View Answer](#)

Answer: c

Explanation: Users see the data in one global schema.

5. Each tablespace in an Oracle database consists of one or more files called

- a) Files
- b) name space
- c) datafiles
- d) PFILE

[View Answer](#)

Answer: c

Explanation: A data file is a computer file which stores data to use by a computer application or system.

6. The management information system (MIS) structure with one main computer system is called a

- a) Hierarchical MIS structure

- b) Distributed MIS structure
- c) Centralized MIS structure
- d) Decentralized MIS structure

[View Answer](#)

Answer: c

Explanation: Structure of MIS may be understood by looking at the physical components of the information system in an organization.

7. A top-to-bottom relationship among the items in a database is established by a

- a) Hierarchical schema
- b) Network schema
- c) Relational schema
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: A hierarchical database model is a data model in which the data is organized into a tree-like structure. The structure allows representing information using parent/child relationships.

8. Choose the RDBMS which supports full fledged client server application development

- a) dBase V
- b) Oracle 7.1
- c) FoxPro 2.1
- d) Ingress

[View Answer](#)

Answer: b

Explanation: RDBMS is Relational Database Management System.

9. One approach to standardization storing of data?

- a) MIS
- b) Structured programming
- c) CODASYL specification
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: CODASYL is an acronym for “Conference on Data Systems Languages”.

10. The highest level in the hierarchy of data organization is called

- a) Data bank
- b) Data base
- c) Data file
- d) Data record

[View Answer](#)

Answer: b

Explanation: Database is a collection of all tables which contains the data in form of fields.

Database Questions and Answers – Organization of Records in Files

1. If a piece of data is stored in two places in the database, then

- a) Storage space is wasted
- b) Changing the data in one spot will cause data inconsistency
- c) It can be more easily accessed
- d) Storage space is wasted & Changing the data in one spot will cause data inconsistency

[View Answer](#)

Answer: d

Explanation: The database is always consistent and so there is no duplication.

2. An audit trail _____

- a) Is used to make backup copies
- b) Is the recorded history of operations performed on a file
- c) Can be used to restore lost information
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: This is more useful for all recovery actions.

3. Large collection of files are called _____

- a) Fields
- b) Records
- c) Database
- d) Sectors

[View Answer](#)

Answer: c

Explanation: The operator tree has a tree like format where the evaluation starts from root of the tree .

4. Which of the following hardware component is the most important to the operation of a database management system?

- a) High resolution video display
- b) Printer
- c) High speed, large capacity disk
- d) Mouse

[View Answer](#)

Answer: c

Explanation: All the data are stored in form of memory in the disk.

5. Which of the following is not true of the traditional approach to information processing

- a) There is common sharing of data among the various applications
- b) It is file oriented
- c) Programs are dependent on the file
- d) It is inflexible

[View Answer](#)

Answer: a

Explanation: All the data are stored in form of memory in the disk.

6. Which of these is not a feature of Hierarchical model?

- a) Organizes the data in tree-like structure
- b) Parent node can have any number of child nodes

- c) Root node does not have any parent
- d) Child node can have any number of parent nodes

[View Answer](#)

Answer: d

Explanation: The data are traversed using several algorithms.

7. Which of these data models is an extension of the relational data model?

- a) Object-oriented data model
- b) Object-relational data model
- c) Semi structured data model
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: All the data are stored in form of memory in the disk.

8. The information about data in a database is called _____

- a) Metadata
- b) Hyper data
- c) Tera data
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Metadata is information about a data.

9. A data dictionary is a special file that contains?

- a) The names of all fields in all files
- b) The data types of all fields in all files
- c) The widths of all fields in all files
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: The data dictionary is structured in tables and views, just like other database data.

10. The DBMS acts as an interface between what two components of an enterprise-class database system?

- a) Database application and the database
- b) Data and the database
- c) The user and the database application
- d) Database application and SQL

[View Answer](#)

Answer: a

Explanation: Database application is the interface with the user to access the database.

Database Questions and Answers – Data-Dictionary Storage

1. A relational database system needs to maintain data about the relations, such as the schema of the relations. This is called

- a) Metadata
- b) Catalog
- c) Log
- d) Dictionary

[View Answer](#)

Answer: a

Explanation: Each side of a platter of a disk has a read–write head that moves across the platter to access different tracks.

2. Relational schemas and other metadata about relations are stored in a structure called the _____

- a) Metadata
- b) Catalog
- c) Log
- d) Data Dictionary

[View Answer](#)

Answer: d

Explanation: Data dictionary is also called as system catalog.

3. _____ is the collection of memory structures and Oracle background processes that operates against an Oracle database.

- a) Database
- b) Instance
- c) Tablespace
- d) Segment

[View Answer](#)

Answer: b

Explanation: Instance is a snapshot of database at any point of time.

4. A _____ is a logical grouping of database objects, usually to facilitate security, performance, or the availability of database objects such as tables and indexes.

- a) Tablespace
- b) Segments
- c) Extents
- d) Blocks

[View Answer](#)

Answer: a

Explanation: A tablespace is a storage location where the actual data underlying database objects can be kept.

5. A tablespace is further broken down into _____

- a) Tablespace
- b) Segments
- c) Extents
- d) Blocks

[View Answer](#)

Answer: b

Explanation: Segment names are used in create table and create index commands to place tables or indexes on specific database devices.

6. _____ is a contiguous group of blocks allocated for use as part of a table, index, and so forth.

- a) Tablespace

b) Segment

c) Extent

d) Block

[View Answer](#)

Answer: c

Explanation: An extent is a set of contiguous blocks allocated in a database.

7. _____ is the smallest unit of allocation in an Oracle database.

a) Database

b) Instance

c) Tablespace

d) Database Block

[View Answer](#)

Answer: d

Explanation: Data block is a form of database space allocation.

8. An Oracle _____ is a set of tables and views that are used as a read-only reference about the database.

a) Database dictionary

b) Dictionary table

c) Data dictionary

d) Dictionary

[View Answer](#)

Answer: c

Explanation: Data dictionary is also called as system catalog.

9. A data dictionary is created when a _____ created.

a) Instance

b) Segment

c) Database

d) Dictionary

[View Answer](#)

Answer: c

Explanation: Data dictionary is also called as system catalog.

10. An Oracle object type has two parts the _____ and _____

a) Instance and body

b) Segment and blocks

c) Specification and body

d) Body and segment

[View Answer](#)

Answer: c

Explanation: Segment names are used in create table and create index commands to place tables or indexes on specific database devices. An extent is a set of contiguous blocks allocated in a database.

Database Questions and Answers – Database Buffer

1. The _____ is that part of main memory available for storage of copies of disk blocks.

- a) Buffer
- b) Catalog
- c) Storage
- d) Secondary storage

[View Answer](#)

Answer: a

Explanation: There is always a copy kept on disk of every block, but the copy on disk may be a version of the block older than the version in the buffer.

2. A major goal of the database system is to minimize the number of block transfers between the disk and memory. This is achieved by

- a) Buffer
- b) Catalog
- c) Storage
- d) Secondary storage

[View Answer](#)

Answer: a

Explanation: There is always a copy kept on disk of every block, but the copy on disk may be a version of the block older than the version in the buffer.

3. The subsystem responsible for the allocation of buffer space is called the _____

- a) Buffer
- b) Buffer manager
- c) Storage
- d) Secondary storage

[View Answer](#)

Answer: b

Explanation: Programs in a database system make requests (that is, calls) on the buffer manager when they need a block from disk.

4. In the buffer where there is no space for another block, the block can be inserted using

- a) Pinned block strategy
- b) Forced output block
- c) Buffer replacement strategy
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Most operating systems use a least recently used (LRU) scheme, in which the block that was referenced least recently is written back to disk and is removed from the buffer.

5. A block that is not allowed to be written back to disk is said to be _____

- a) Pinned
- b) Forced
- c) Buffer
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Although many operating systems do not support pinned blocks, such a feature is essential for a database system that is resilient to crashes.

6. There are situations in which it is necessary to write back the block to disk, even though the buffer space that it occupies is not needed. This write is called the

- a) Pinned block strategy
- b) Forced output block
- c) Buffer replacement strategy
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: The main-memory contents and thus buffer contents are lost in a crash, whereas data on disk usually survive a crash.

7. The frequently used buffer replacement strategy is

- a) Most recently used
- b) Least recently used
- c) Longest block
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: If a block must be replaced, the least recently referenced block is replaced.

8. In case the buffer manager do not write the blocks properly then the buffer manager uses

- a) Replacement strategy
- b) Forced strategy
- c) Crash recovery system
- d) Both Replacement and Forced strategy

[View Answer](#)

Answer: c

Explanation: The crash-recovery subsystem imposes stringent constraints on block replacement.

9. The technique where the blocks which have been used are replaced is called

- a) Replacement strategy
- b) Forced strategy
- c) Crash recovery system
- d) Most recently used

[View Answer](#)

Answer: d

Explanation: The optimal strategy for block replacement is the most recently used (MRU) strategy.

10. _____ frees the space occupied by a block as soon as the final tuple of that block has been processed.

- a) Replacement strategy
- b) Forced strategy
- c) Toss immediate strategy
- d) Most recently used

[View Answer](#)

Answer: c

Explanation: The optimal strategy for block replacement is the most recently used (MRU) strategy.

Database Questions and Answers – Ordered Indices

1. In ordered indices the file containing the records is sequentially ordered, a _____ is an index whose search key also defines the sequential order of the file.

- a) Clustered index
- b) Structured index
- c) Unstructured index
- d) Nonclustered index

[View Answer](#)

Answer: a

Explanation: Clustering index are also called primary indices; the term primary index may appear to denote an index on a primary key, but such indices can in fact be built on any search key.

2. Indices whose search key specifies an order different from the sequential order of the file are called _____ indices.

- a) Nonclustered
- b) Secondary
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Nonclustering index is also called secondary indices.

3. An _____ consists of a search-key value and pointers to one or more records with that value as their search-key value.

- a) Index entry
- b) Index hash
- c) Index cluster
- d) Index map

[View Answer](#)

Answer: a

Explanation: The pointer to a record consists of the identifier of a disk block and an offset within the disk block to identify the record within the block.

4. In a _____ clustering index, the index record contains the search-key value and a pointer to the first data record with that search-key value and the rest of the records will be in the sequential pointers.

- a) Dense
- b) Sparse
- c) Straight
- d) Continuous

[View Answer](#)

Answer: a

Explanation: In a dense nonclustering index, the index must store a list of pointers to all records with the same search-key value.

5. In a _____ index, an index entry appears for only some of the search-key values.

- a) Dense
- b) Sparse
- c) Straight
- d) Continuous

[View Answer](#)

Answer: a

Explanation: Sparse indices can be used only if the relation is stored in sorted order of the search key, that is if the index is a clustering index.

6. Incase the indices values are larger, index is created for these values of the index. This is called

- a) Pointed index
- b) Sequential index
- c) Multilevel index
- d) Multiple index

[View Answer](#)

Answer: c

Explanation: Indices with two or more levels are called multilevel indices.

7. A search key containing more than one attribute is referred to as a _____ search key.

- a) Simple
- b) Composite
- c) Compound
- d) Secondary

[View Answer](#)

Answer: b

Explanation: The structure of the index is the same as that of any other index, the only difference being that the search key is not a single attribute, but rather is a list of attributes.

8. In B+ tree the node which points to another node is called

- a) Leaf node
- b) External node
- c) Final node
- d) Internal node

[View Answer](#)

Answer: d

Explanation: Nonleaf nodes are also referred to as internal nodes.

9. Insertion of a large number of entries at a time into an index is referred to as _____ of the index.

- a) Loading
- b) Bulk insertion
- c) Bulk loading
- d) Increase insertion

[View Answer](#)

Answer: c

Explanation: Bulk loading is used to improve efficiency and scalability.

10. While inserting the record into the index, if the search-key value does not appear in the index.

- a) The system adds a pointer to the new record in the index entry
- b) The system places the record being inserted after the other records with the same search-key values
- c) The system inserts an index entry with the search-key value in the index at the appropriate position
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: If the index entry stores pointers to all records with the same search key value, the system adds a pointer to the new record in the index entry.

Database Questions and Answers – Hashing techniques

1. If h is any hashing function and is used to hash n keys into a table of size m , where $n \leq m$, the expected number of collisions involving a particular key x is :

- a) Less than 1
- b) Less than n
- c) Less than m
- d) Less than $n/2$

[View Answer](#)

Answer: a

Explanation: Hashing is also a method of sorting key values in a database table in an efficient manner.

2. A technique for direct search is

- a) Binary Search
- b) Linear Search
- c) Tree Search
- d) Hashing

[View Answer](#)

Answer: d

Explanation: Hashing is one way to enable security during the process of message transmission when the message is intended for a particular recipient only.

3. The searching technique that takes $O(1)$ time to find a data is

- a) Linear Search
- b) Binary Search
- c) Hashing
- d) Tree Search

[View Answer](#)

Answer: c

Explanation: A formula generates the hash, which helps to protect the security of the transmission from unauthorized users.

4. The goal of hashing is to produce a search that takes

- a) $O(1)$ time
- b) $O(n^2)$ time
- c) $O(\log n)$ time
- d) $O(n \log n)$ time

[View Answer](#)

Answer: a

Explanation: Time complexity is given by the big oh notation.

5. Consider a hash table of size seven, with starting index zero, and a hash function $(3x + 4) \bmod 7$. Assuming the hash table is initially empty, which of the following is the contents of the table when the sequence 1, 3, 8, 10 is inserted into the table using closed hashing? Note that ‘_’ denotes an empty location in the table.

- a) 8, _, _, _, _, _, 10
- b) 1, 8, 10, _, _, _, 3
- c) 1, _, _, _, _, 3
- d) 1, 10, 8, _, _, _, 3

[View Answer](#)

Answer: b

Explanation: A formula generates the hash, which helps to protect the security of the transmission from unauthorized users.

6. A hash table can store a maximum of 10 records, currently there are records in location 1, 3,4,7,8,9,10. The probability of a new record going into location 2, with hash functions resolving collisions by linear probing is

- a) 0.1
- b) 0.6
- c) 0.2
- d) 0.5

[View Answer](#)

Answer: b

Explanation: Hashing is used to index and retrieve items in a database because it is easier to find the item using the shortened hashed key than using the original value.

7. Key value pairs is usually seen in

- a) Hash tables
- b) Heaps
- c) Both Hash tables and Heaps
- d) Skip list

[View Answer](#)

Answer: a

Explanation: Hashing is used to index and retrieve items in a database because it is easier to find the item using the shortened hashed key than using the original value.

8. What is the best definition of a collision in a hash table?

- a) Two entries are identical except for their keys
- b) Two entries with different data have the exact same key
- c) Two entries with different keys have the same exact hash value
- d) Two entries with the exact same key have different hash values

[View Answer](#)

Answer: a

Explanation: This level is the root of the tree.

9. Which of the following scenarios leads to linear running time for a random search hit in a linear-probing hash table?

- a) All keys hash to same index
- b) All keys hash to different indices
- c) All keys hash to an even-numbered index
- d) All keys hash to different even-numbered indices

[View Answer](#)

Answer: a

Explanation: If all keys hash to the same location then the i -th inserted key would need i lookups to be found. The probability of looking up i -th key is $1/n$ (since it's random). If you know some probability it's trivial to show that such lookups have linear time.

10. Breadth First Search is used in

- a) Binary trees
- b) Stacks
- c) Graphs
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Hashing is used to index and retrieve items in a database because it is easier to find the item using the shortened hashed key than using the original value.

Database Questions and Answers – Ordered Indexing and Hashing

1. A(n) _____ can be used to preserve the integrity of a document or a message.

- a) Message digest
- b) Message summary
- c) Encrypted message
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Encryption algorithms are used to keep the contents safe.

2. A hash function must meet _____ criteria.

- a) Two
- b) Three
- c) Four
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Only if the criteria is fulfilled the values are hashed.

3. What is the main limitation of Hierarchical Databases?

- a) Limited capacity (unable to hold much data)
- b) Limited flexibility in accessing data
- c) Overhead associated with maintaining indexes
- d) The performance of the database is poor

[View Answer](#)

Answer: b

Explanation: In this, the data items are placed in a tree like hierarchical structure.

4. The property (or set of properties) that uniquely defines each row in a table is called the:

- a) Identifier
- b) Index
- c) Primary key
- d) Symmetric key

[View Answer](#)

Answer: c

Explanation: Primary is used to uniquely identify the tuples.

5. The separation of the data definition from the program is known as:

- a) Data dictionary
- b) Data independence
- c) Data integrity
- d) Referential integrity

[View Answer](#)

Answer: b

Explanation: Data dictionary is the place where the meaning of the data are organized.

6. In the client / server model, the database:

- a) Is downloaded to the client upon request
- b) Is shared by both the client and server

- c) Resides on the client side
- d) Resides on the server side

[View Answer](#)

Answer: d

Explanation: The server has all the database information and the client access it.

7. The traditional storage of data that is organized by customer, stored in separate folders in filing cabinets is an example of what type of 'database' system?

- a) Hierarchical
- b) Network
- c) Object oriented
- d) Relational

[View Answer](#)

Answer: a

Explanation: Hierarchy is based on Parent-Child Relationship. Parent-Child Relationship Type is basically 1:N relationship.

8. The database design that consists of multiple tables that are linked together through matching data stored in each table is called

- a) Hierarchical database
- b) Network database
- c) Object oriented database
- d) Relational database

[View Answer](#)

Answer: d

Explanation: A relational database is a collection of data items organized as a set of formally described tables from which data can be accessed or reassembled.

9. The association role defines:

- a) How tables are related in the database
- b) The relationship between the class diagram and the tables in the database
- c) The tables that each attribute is contained
- d) Which attribute is the table's primary key

[View Answer](#)

Answer: a

Explanation: The tables are always related in the database to form consistency.

10. The purpose of an N-Ary association is:

- a) To capture a parent-child relationship
- b) To deal with one to many relationships
- c) To deal with relationships that involve more than two tables
- d) To represent an inheritance relationship

[View Answer](#)

Answer: c

Explanation: This is binary n-ary association meaning more than two classes are involved in the relationship.

Database Questions & Answers – Bitmap Indices

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

1. Bitmap indices are a specialized type of index designed for easy querying on _____

- a) Bit values
- b) Binary digits
- c) Multiple keys
- d) Single keys

[View Answer](#)

Answer: c

Explanation: Each bitmap index is built on a single key.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

2. A _____ on the attribute A of relation r consists of one bitmap for each value that A can take.

- a) Bitmap index
- b) Bitmap
- c) Index
- d) Array

[View Answer](#)

Answer: a

Explanation: A bitmap is simply an array of bits.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

In this selection, we fetch the bitmaps for gender value f and the bitmap for income level value L2, and perform an _____ of the two bitmaps.

- a) Union
- b) Addition
- c) Combination
- d) Intersection

[View Answer](#)

Answer: d

Explanation: We compute a new bitmap where bit i has value 1 if the ith bit of the two bitmaps are both 1, and has a value 0 otherwise.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

4. To identify the deleted records we use the _____

- a) Existence bitmap

- b) Current bitmap
- c) Final bitmap
- d) Deleted bitmap

[View Answer](#)

Answer: a

Explanation: The bitmaps which are deleted are denoted by 0.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

5. Bitmaps can be used as a compressed storage mechanism at the leaf nodes of _____ for those values that occur very frequently.

- a) B-trees
- b) B+-trees
- c) Bit trees
- d) Both B-trees and B+-trees

[View Answer](#)

Answer: b

Explanation: Bitmaps are combined and stored in a B+ tree.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

6. Bitmaps can be combined with regular B+-tree indices for relations where a few attribute values are extremely common, and other values also occur, but much less frequently.

- a) Bitmap, B-tree
- b) Bitmap, B+tree
- c) B-tree, Bitmap
- d) B+tree, Bitmap

[View Answer](#)

Answer: b

Explanation: Bitmaps are combined and stored in a B+ tree.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

7. In a B+-tree index _____ for each value, we would normally maintain a list of all records with that value for the indexed attribute.

- a) Leaf
- b) Node
- c) Root
- d) Link

[View Answer](#)

Answer: a

Explanation: Bitmaps are combined and stored in a B+ tree.

3.

```
SELECT *
```

```
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

8. A tablespace is further broken down into _____

- a) Tablespace
- b) Segments
- c) Extents
- d) Blocks

[View Answer](#)

Answer: b

Explanation: Segment names are used in create table and create index commands to place tables or indexes on specific database devices.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

9. In ordered indices the file containing the records is sequentially ordered, a _____ is an index whose search key also defines the sequential order of the file.

- a) Clustered index
- b) Structured index
- c) Unstructured index
- d) Nonclustered index

[View Answer](#)

Answer: a

Explanation: Clustering index are also called primary indices; the term primary index may appear to denote an index on a primary key, but such indices can in fact be built on any search key.

3.

```
SELECT *
FROM r
WHERE gender = 'f' AND income level = 'L2';
```

10. Indices whose search key specifies an order different from the sequential order of the file are called _____ indices.

- a) Nonclustered
- b) Secondary
- c) All of the mentioned
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Nonclustering index are also called secondary indices.

Database Questions and Answers – Index Definition in SQL

1. What is the purpose of the index in sql server?

- a) To enhance the query performance
- b) To provide an index to a record
- c) To perform fast searches
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: A database index is a data structure that improves the speed of data retrieval operations on a database table at the cost of additional writes.

2. How many types of indexes are there in sql server?

- a) 1
- b) 2
- c) 3
- d) 4

[View Answer](#)

Answer: b

Explanation: They are clustered index and non clustered index.

3. How non clustered index point to the data?

- a) It never points to anything
- b) It points to a data row
- c) It is used for pointing data rows containing key values
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Nonclustered indexes have a structure separate from the data rows. A nonclustered index contains the nonclustered index key values and each key value entry has a pointer to the data row that contains the key value.

4. Which one is true about clustered index?

- a) Clustered index is not associated with table
- b) Clustered index is built by default on unique key columns
- c) Clustered index is not built on unique key columns
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Nonclustered indexes have a structure separate from the data rows. A nonclustered index contains the nonclustered index key values and each key value entry has a pointer to the data row that contains the key value.

5. What is true about indexes?

- a) Indexes enhance the performance even if the table is updated frequently
- b) It makes harder for sql server engines to work to work on index which have large keys
- c) It doesn't make harder for sql server engines to work to work on index which have large keys
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Indexes tend to improve the performance.

6. Does index take space in the disk?

- a) It stores memory as and when required
- b) Yes, Indexes are stored on disk
- c) Indexes are never stored on disk
- d) Indexes take no space

[View Answer](#)

Answer: b

Explanation: Indexes take memory slots which are located on the disk.

7. What are composite indexes?

- a) Are those which are composed by database for its internal use
- b) A composite index is a combination of index on 2 or more columns
- c) Composite index can never be created
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: A composite index is an index on two or more columns of a table.

8. If an index is _____ the metadata and statistics continue to exists

- a) Disabling
- b) Dropping
- c) Altering
- d) Both Disabling and Dropping

[View Answer](#)

Answer: a

Explanation: A database index is a data structure that improves the speed of data retrieval operations on a database table at the cost of additional writes.

9. In _____ index instead of storing all the columns for a record together, each column is stored separately with all other rows in an index.

- a) Clustered
- b) Column store
- c) Non clustered
- d) Row store

[View Answer](#)

Answer: b

Explanation: A database index is a data structure that improves the speed of data retrieval operations on a database table at the cost of additional writes.

10. A _____ index is the one which satisfies all the columns requested in the query without performing further lookup into the clustered index.

- a) Clustered
- b) Non Clustered
- c) Covering
- d) B-Tree

[View Answer](#)

Answer: c

Explanation: A covered query is a query where all the columns in the query's result set are pulled from non-clustered indexes.

Database Questions and Answers – Query Processing

1. A collection of data designed to be used by different people is called a/an

- a) Organization
- b) Database
- c) Relationship
- d) Schema

[View Answer](#)

Answer: b

Explanation: Database is a collection of related tables.

2. Which of the following is the oldest database model?

- a) Relational
- b) Deductive
- c) Physical
- d) Network

[View Answer](#)

Answer: d

Explanation: The network model is a database model conceived as a flexible way of representing objects and their relationships.

3. Which of the following schemas does define a view or views of the database for particular users?

- a) Internal schema
- b) Conceptual schema
- c) Physical schema
- d) External schema

[View Answer](#)

Answer: d

Explanation: An externally-defined schema can provide access to tables that are managed on any PostgreSQL, Microsoft SQL Server, SAS, Oracle, or MySQL database.

4. Which of the following is an attribute that can uniquely identify a row in a table?

- a) Secondary key
- b) Candidate key
- c) Foreign key
- d) Alternate key

[View Answer](#)

Answer: b

Explanation: A Candidate Key can be any column or a combination of columns that can qualify as unique key in database.

5. Which of the following are the process of selecting the data storage and data access characteristics of the database?

- a) Logical database design
- b) Physical database design
- c) Testing and performance tuning
- d) Evaluation and selecting

[View Answer](#)

Answer: b

Explanation: The physical design of the database optimizes performance while ensuring data integrity by avoiding unnecessary data redundancies.

6. Which of the following terms does refer to the correctness and completeness of the data in a database?

- a) Data security
- b) Data constraint
- c) Data independence
- d) Data integrity

[View Answer](#)

Answer: d

Explanation: ACID property is satisfied by transaction in database.

7. The relationship between DEPARTMENT and EMPLOYEE is a

- a) One-to-one relationship
- b) One-to-many relationship
- c) Many-to-many relationship
- d) Many-to-one relationship

[View Answer](#)

Answer: b

Explanation: One entity department is related to several employees.

8. A table can be logically connected to another table by defining a

- a) Super key
- b) Candidate key
- c) Primary key
- d) Unique key

[View Answer](#)

Answer: c

Explanation: A superkey is a combination of attributes that can be uniquely used to identify a database record.

9. If the state of the database no longer reflects a real state of the world that the database is supposed to capture, then such a state is called

- a) Consistent state
- b) Parallel state
- c) Durable state
- d) Inconsistent state

[View Answer](#)

Answer: d

Explanation: SQL data consistency is that whenever a transaction is performed, it sees a consistent database.

10. Ensuring isolation property is the responsibility of the

- a) Recovery-management component of the DBMS
- b) Concurrency-control component of the DBMS
- c) Transaction-management component of the DBMS
- d) Buffer management component in DBMS

[View Answer](#)

Answer: b

Explanation: Concurrency control ensures that correct results for concurrent operations are generated while getting those results as quickly as possible.

Database Questions and Answers – Selection Operation

1. In query processing, the _____ is the lowest-level operator to access data.

- a) Index Search
- b) Linear search
- c) File scan
- d) Access paths

[View Answer](#)

Answer: c

Explanation: File scans are search algorithms that locate and retrieve records that fulfill a selection condition.

2. In a _____ the system scans each file block and tests all records to see whether they satisfy the selection condition.

- a) Index Search
- b) Linear search
- c) File scan
- d) Access paths

[View Answer](#)

Answer: b

Explanation: An initial seek is required to access the first block of the file.

3. Index structures are referred to as _____ since they provide a path through which data can be located and accessed.

- a) Index Search
- b) Linear search
- c) File scan
- d) Access paths

[View Answer](#)

Answer: d

Explanation: A primary index is an index that allows the records of a file to be read in an order that corresponds to the physical order in the file.

4. Search algorithms that use an index are referred to as

- a) Index Search
- b) Linear search
- c) File scan
- d) Access paths

[View Answer](#)

Answer: a

Explanation: Selection predicates are used to guide in the choice of the index to use in processing the query.

5. Which algorithm uses equality comparison on a key attribute with a primary index to retrieve a single record that satisfies the corresponding equality condition.

- a) A2
- b) A4
- c) A5
- d) A6

[View Answer](#)

Answer: a

Explanation: A2 – primary index, equality on key.

6. The strategy can retrieve a single record if the equality condition is on a key; multiple records may be retrieved if the indexing field is not a key is

- a) A2
- b) A4
- c) A5
- d) A6

[View Answer](#)

Answer: b

Explanation: A4 – Secondary index, equality.

7. The algorithm that uses a secondary ordered index to guide retrieval for comparison conditions involving $<$, \leq , \geq , or $>$ is

- a) A2
- b) A4
- c) A5
- d) A6

[View Answer](#)

Answer: d

Explanation: A6 – Secondary index, comparison.

8. The ___ algorithm scans each index for pointers to tuples that satisfy an individual condition.

- a) A2
- b) A4
- c) A9
- d) A6

[View Answer](#)

Answer: c

Explanation: A9 – Conjunctive selection by an intersection of identifiers.

9. If access paths are available on all the conditions of a disjunctive selection, each index is scanned for pointers to tuples that satisfy the individual condition. This is satisfied by

- a) A10
- b) A7
- c) A9
- d) A6

[View Answer](#)

Answer: a

Explanation: A10 – Disjunctive selection by union of identifiers.

10. Conjunctive selection using one index. This is

- a) A10
- b) A7
- c) A9
- d) A6

[View Answer](#)

Answer: b

Explanation: To reduce the cost of A7 we choose a i and one of algorithms A1 through A6 for which the combination results in the least cost for $i(r)$. The cost of algorithm A7 is given by the cost of the chosen algorithm.

Database Questions and Answers – Sorting

1. Two main measures for the efficiency of an algorithm are

- a) Processor and memory
- b) Complexity and capacity
- c) Time and space
- d) Data and space

[View Answer](#)

Answer: c

Explanation: Depending on the time and space complexity only the algorithm for sorting will be chosen.

2. The time factor when determining the efficiency of an algorithm is measured by

- a) Counting microseconds
- b) Counting the number of key operations
- c) Counting the number of statements
- d) Counting the kilobytes of algorithm

[View Answer](#)

Answer: b

Explanation: The operations taking place with the time and space is counted.

3. The space factor when determining the efficiency of an algorithm is measured by

- a) Counting the maximum memory needed by the algorithm
- b) Counting the minimum memory needed by the algorithm
- c) Counting the average memory needed by the algorithm
- d) Counting the maximum disk space needed by the algorithm

[View Answer](#)

Answer: a

Explanation: Time complexity maintains the maximum time needed.

4. Which of the following case does not exist in complexity theory

- a) Best case
- b) Worst case
- c) Average case
- d) Null case

[View Answer](#)

Answer: d

Explanation: Null case cannot be counted as the factor for complexity.

5. The Worst case occur in linear search algorithm when

- a) Item is somewhere in the middle of the array
- b) Item is not in the array at all
- c) Item is the last element in the array
- d) Item is the last element in the array or is not there at all

[View Answer](#)

Answer: d

Explanation: Algorithmic complexity is concerned about how fast or slow particular algorithm performs.

6. The Average case occur in linear search algorithm

- a) When Item is somewhere in the middle of the array
- b) When Item is not in the array at all

- c) When Item is the last element in the array
- d) When Item is the last element in the array or is not there at all

[View Answer](#)

Answer: a

Explanation: Algorithmic complexity is concerned about how fast or slow particular algorithm performs.

7. The complexity of the average case of an algorithm is

- a) Much more complicated to analyze than that of worst case
- b) Much more simpler to analyze than that of worst case
- c) Sometimes more complicated and some other times simpler than that of worst case
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Algorithmic complexity is concerned about how fast or slow particular algorithm performs.

8. The complexity of a linear search algorithm is

- a) $O(n)$
- b) $O(\log n)$
- c) $O(n^2)$
- d) $O(n \log n)$

[View Answer](#)

Answer: a

Explanation: It refers to n values complexity in the algorithm which can be reduced by choosing the other algorithms.

9. The complexity of Binary search algorithm is

- a) $O(n)$
- b) $O(\log n)$
- c) $O(n^2)$
- d) $O(n \log n)$

[View Answer](#)

Answer: b

Explanation: This shows that it has a standard complexity in addressing.

10. The complexity of Bubble sort algorithm is

- a) $O(n)$
- b) $O(\log n)$
- c) $O(n^2)$
- d) $O(n \log n)$

[View Answer](#)

Answer: c

Explanation: Bubble sort, is a simple sorting algorithm that works by repeatedly stepping through the list to be sorted, comparing each pair of adjacent items and swapping them if they are in the wrong order.

Database Questions and Answers – Join Operations

1. A _____ is a query that retrieves rows from more than one table or view:

- a) Start
- b) End
- c) Join
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: An SQL join clause combines records from two or more tables in a database. It creates a set that can be saved as a table or used as it is. A JOIN is a means for combining fields from two tables by using values common to each.

2. A condition is referred to as _____

- a) Join in SQL
- b) Join condition
- c) Join in SQL & Condition
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: An SQL join clause combines records from two or more tables in a database. It creates a set that can be saved as a table or used as it is. A JOIN is a means for combining fields from two tables by using values common to each.

3. Which oracle is the join condition is specified using the WHERE clause:

- a) Oracle 9i
- b) Oracle 8i
- c) Pre-oracle 9i
- d) Pre-oracle 8i

[View Answer](#)

Answer: c

Explanation: Oracle 9i is a version of the Oracle Database. The i stands for “Internet” to indicate that 9i is “Internet ready”.

4. How many join types in join condition:

- a) 2
- b) 3
- c) 4
- d) 5

[View Answer](#)

Answer: d

Explanation: INNER JOIN, LEFT JOIN, RIGHT JOIN, FULL JOIN, EQUIJOIN.

5. Which are the join types in join condition:

- a) Cross join
- b) Natural join
- c) Join with USING clause
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: INNER JOIN, LEFT JOIN, RIGHT JOIN, FULL JOIN, EQUIJOIN are the types of joins.

6. Which product is returned in a join query have no join condition:

- a) Equijoins
- b) Cartesian
- c) Both Equijoins and Cartesian
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: A Cartesian coordinate system is a coordinate system that specifies each point uniquely in a plane by a pair of numerical coordinates.

7. Which is a join condition contains an equality operator:

- a) Equijoins
- b) Cartesian
- c) Both Equijoins and Cartesian
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: An equi-join is a specific type of comparator-based join, that uses only equality comparisons in the join-predicate.

8. Which join refers to join records from the write table that have no matching key in the left table are include in the result set:

- a) Left outer join
- b) Right outer join
- c) Full outer join
- d) Half outer join

[View Answer](#)

Answer: b

Explanation: A right outer join will return all the rows that an inner join returns plus one row for each of the other rows in the second table that did not have a match in the first table. It is the same as a left outer join with the tables specified in the opposite order.

9. Which operation are allowed in a join view:

- a) UPDATE
- b) INSERT
- c) DELETE
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: The DELETE statement is used to delete rows in a table. The UPDATE statement is used to update existing records in a table. The INSERT INTO statement is used to insert new records in a table.

10. Which view that contains more than one table in the top-level FROM clause of the SELECT statement:

- a) Join view
- b) Datable join view
- c) Updatable join view
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: The DELETE statement is used to delete rows in a table. The UPDATE statement is used to update existing records in a table. The INSERT INTO statement is used to insert new records in a table.

Database Questions and Answers – Evaluation of Expressions

1. Pictorial representation of an expression is called

- a) Expression tree
- b) Operator tree
- c) Expression flow
- d) Expression chart

[View Answer](#)

Answer: b

Explanation: The operator tree has a tree like format where the evaluation starts from root of the tree.

2. The results of each intermediate operation are created and then are used for evaluation of the next-level operations. This is called

- a) Materialized evaluation
- b) Expression evaluation
- c) Tree evaluation
- d) Tree materialization

[View Answer](#)

Answer: a

Explanation: The cost of a materialized evaluation is not simply the sum of the costs of the operations involved.

3. _____ allows the algorithm to execute more quickly by performing CPU activity in parallel with I/O activity.

- a) Buffering
- b) Double buffering
- c) Multiple buffering
- d) Double reading

[View Answer](#)

Answer: a

Explanation: Double buffering using two buffers, with one continuing execution of the algorithm while the other is being written out.

4. Pipelines can be executed in

- a) 4
- b) 3
- c) 2
- d) 5

[View Answer](#)

Answer: c

Explanation: Demand driven and producer driven pipelines are the two ways.

5. In a _____ the system makes repeated requests for tuples from the operation at the top of the pipeline.

- a) Demand-driven pipeline
- b) Producer-driven pipeline
- c) Demand pipeline
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Each time that an operation receives a request for tuples, it computes the next tuple (or tuples) to be returned, and then returns that tuple.

6. In a _____ operations do not wait for requests to produce tuples, but instead generate the tuples eagerly.

- a) Demand-driven pipeline

- b) Producer-driven pipeline
- c) Demand pipeline
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: Each operation in a producer-driven pipeline is modeled as a separate process or thread within the system that takes a stream of tuples from its pipelined inputs and generates a stream of tuples for its output.

7. Each operation in a demand-driven pipeline can be implemented as an _____ that provides the following functions: open(), next(), and close().

- a) Demand
- b) Pipeline
- c) Iterator
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: After a call to open(), each call to next() returns the next output tuple of the operation.

8. The iterator maintains the _____ of its execution in between calls so that successive next() requests receive successive result tuples.

- a) State
- b) Transition
- c) Rate
- d) Block

[View Answer](#)

Answer: a

Explanation: The function close() tells an iterator that no more tuples are required.

9. Tuples are generated _____ in producer-driven pipelining, they are generated _____ on demand, in demand-driven pipelining.

- a) Lazily, Eagerly
- b) Eagerly, Lazily
- c) Slowly, Eagerly
- d) Eagerly, Slowly

[View Answer](#)

Answer: b

Explanation: Producer-driven pipelining is very useful in parallel processing systems.

10. When two inputs that we desire to pipeline into the join are not already sorted it is the _____ technique.

- a) Hash join
- b) Buffer join
- c) double-pipelined hash join
- d) double-pipelined join

[View Answer](#)

Answer: d

Explanation: When hash indices are used on tuples, the resultant algorithm is called the double-pipelined hash-join technique.

Database Questions and Answers – Transformation of Relational Expressions

Assume {Author, Title} is the key for both schemes. Which of the following statements is true?

- a) Both Book and Collection are in BCNF
- b) Both Book and Collection are in 3NF only
- c) Book is in 2NF and Collection is in 3NF
- d) Both Book and Collection are in 2NF only

[View Answer](#)

Answer: c

Explanation: The relation Collection is in BCNF: Its given that {Author, Title} is the key and there is only one functional dependency (FD) applicable to the relation Collection {i.e. Title Author \rightarrow Catalog_no}.

The relation schema R is

- a) in BCNF
- b) in 3NF, but not in BCNF
- c) in 2NF, but not in 3NF
- d) not in 2NF

[View Answer](#)

Answer: d

Explanation: From the closure set of attributes we can see that the key for the relation is AB. The FD B \rightarrow G is a partial dependency, hence it is not in 2NF.

3. Which of the following is/are false for RAW mode of FOR XML?

- a) XMLSCHEMA option does not returns an in-line XSD schema
- b) BINARY BASE32 returns the binary data in base32-encoded format
- c) Each row in the query result is transformed into an XML element
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: XML was designed to transport and store data.

4. _____ refers to the ability of the system to recover committed transaction updates if either the system or the storage media fails.

- a) Isolation
- b) Atomicity
- c) Consistency
- d) Durability

[View Answer](#)

Answer: d

Explanation: In database systems, durability is the ACID property which guarantees that transactions that have committed will survive permanently.

5. Which utilities can we use to export data from sql server to a text file?

- a) DTS export wizard
- b) BCP
- c) ISQL
- d) DTS export wizard and BCP

[View Answer](#)

Answer: d

Explanation: The bcp utility bulk copies data between an instance of Microsoft SQL Server and a data file in a user-specified format.

6. You have a column that will only contain values from 0 to 256. What is the most economical data type to use for the column?

- a) TINYINT
- b) SMALLINT
- c) INT
- d) DECIMAL(1)

[View Answer](#)

Answer: b

Explanation: The bcp utility bulk copies data between an instance of Microsoft SQL Server and a data file in a user-specified format.

7. Problems occurs if we don't implement a proper locking strategy

- a) Dirty reads
- b) Phantom reads
- c) Lost updates
- d) Unrepeatable reads

[View Answer](#)

Answer: b

Explanation: Phantom reads occur when an insert or delete action is performed against a row that belongs to a range of rows being read by a transaction.

8. Which of the following fixed database roles can add or remove user IDs?

- a) db_accessadmin
- b) db_securityadmin
- c) db_setupadmin
- d) db_sysadmin

[View Answer](#)

Answer: a

Explanation: The db_accessadmin role manages security, but handles access to the database, as the name implies.

9. By default sql server has _____ isolation level

- a) READ COMMITTED
- b) READ UNCOMMITTED
- c) SERIALIZABLE
- d) REPEATABLE READ

[View Answer](#)

Answer: a

Explanation: READ UNCOMMITTED is the most optimistic concurrency isolation option available in SQL Server.

10. Which of the following pair of regular expression are not equivalent?

- a) $1(01)^*$ and $(10)^*1$
- b) $x(xx)^*$ and $(xx)^*x$
- c) $(ab)^*$ and a^*b^*
- d) x^+ and x^*x^+

[View Answer](#)

Answer: c

Explanation: $(ab)^* = (a^*b^*)^*$.

Database Questions and Answers – Estimating Statistics of Expression Results

1. Which feature converts row data to a column for better analytical view?

- a) Views
- b) Join
- c) Pivot
- d) Trigger

[View Answer](#)

Answer: c

Explanation: Pivot table is very powerful and very easy to use.

2. Which of the following statements is/are not true for SQL profiler?

- a) Enables you to monitor events
- b) Check if rows are being inserted properly
- c) Check the performance of a stored procedure
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Stored procedures are like functions which do not return values.

3. Which global variables can be used to determine if a transaction is still open?

- a) @@NESTLEVEL
- b) @@FETCH_STATUS
- c) @@TRANCOUNT
- d) @@CONNECTIONS

[View Answer](#)

Answer: c

Explanation: PRINT @@TRANCOUNT — The BEGIN TRAN statement will increment the — transaction count by 1.

4. Which statement is used to define a cursor?

- a) OPEN
- b) FETCH
- c) DECLARE CURSOR
- d) @@FETCH_STATUS

[View Answer](#)

Answer: c

Explanation: A database cursor is a control structure that enables traversal over the records in a database.

5. What is the default “SORT” order for a SQL?

- a) Ascending
- b) Descending
- c) As specified by the user
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Default is ascending order.

6. Capabilities of RAISERROR

- a) It can be logged in the error log
- b) It can print a message to the application

- c) It can assign an error number, state and severity
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: A relational database table is often described as “normalized” if it is in the Third Normal Form because most of the 3NF tables are free of insertion, update, and deletion anomalies.

7. How inserting data through stored procedure do reduces network traffic and increase database performance?

- a) Stored procedure can accept parameter
- b) Permission check is not required
- c) The execution plan is stored in the cache after it was executed the first time
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: A relational database table is often described as “normalized” if it is in the Third Normal Form because most of the 3NF tables are free of insertion, update, and deletion anomalies.

8. Stored procedures are safe from SQL injection attacks

- a) True
- b) False
- c) Depends on the result
- d) Always safe

[View Answer](#)

Answer: a

Explanation: Injection attack is not possible in SP.

9. Which of the following connection type supports application role permissions and password encryption?

- a) OLE DB
- b) DBLib
- c) ODBC
- d) OLE DB and ODBC

[View Answer](#)

Answer: d

Explanation: Open Database Connectivity (ODBC) is Microsoft’s strategic interface for accessing data in a heterogeneous environment of relational.

10. Cursor that reflects the changes made to the database table even after the result set is returned

- a) Static
- b) Dynamic
- c) FORWARD_ONLY
- d) Keyset

[View Answer](#)

Answer: b

Explanation: A database cursor is a control structure that enables traversal over the records in a database.

Database Questions and Answers – Materialized Views

1. Which normal form is considered adequate for normal relational database design?

- a) 2NF
- b) 5NF
- c) 4NF
- d) 3NF

[View Answer](#)

Answer: d

Explanation: A relational database table is often described as “normalized” if it is in the Third Normal Form because most of the 3NF tables are free of insertion, update, and deletion anomalies.

2. Consider a schema R(A, B, C, D) and functional dependencies A \rightarrow B and C \rightarrow D. Then the decomposition of R into R1 (A, B) and R2(C, D) is

- a) dependency preserving and lossless join
- b) lossless join but not dependency preserving
- c) dependency preserving but not lossless join
- d) not dependency preserving and not lossless join

[View Answer](#)

Answer: d

Explanation: While decomposing a relational table we must verify the following properties:

- i) Dependency Preserving Property
- ii) Lossless-Join Property.

3. Relation R with an associated set of functional dependencies, F, is decomposed into BCNF. The redundancy (arising out of functional dependencies) in the resulting set of relations is

- a) Zero
- b) More than zero but less than that of an equivalent 3NF decomposition
- c) Proportional to the size of F⁺
- d) Indeterminate

[View Answer](#)

Answer: b

Explanation: Redundancy in BCNF is low when compared to 3NF.

4. Which one of the following statements about normal forms is FALSE?

- a) BCNF is stricter than 3NF
- b) Lossless, dependency-preserving decomposition into 3NF is always possible
- c) Lossless, dependency-preserving decomposition into BCNF is always possible
- d) Any relation with two attributes is in BCNF

[View Answer](#)

Answer: c

Explanation: Achieving Lossless and dependency-preserving decomposition property into BCNF is difficult.

in terms of normalization, this table is in

- a) 1NF
- b) 2NF
- c) 3NF
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Since the primary key is not given we have to derive the primary key of the table. Using the closure set of attributes we get the primary key as (F1,F2). From functional dependencies, “F1->F3, F2->F4”, we can see that there is partial functional dependency therefore it is not in 1NF. Hence the table is in 1NF.

6. Which of the following is TRUE?

- a) Every relation in 2NF is also in BCNF
- b) A relation R is in 3NF if every non-prime attribute of R is fully functionally dependent on every key of R
- c) Every relation in BCNF is also in 3NF
- d) No relation can be in both BCNF and 3NF

[View Answer](#)

Answer: c

Explanation: A relational database table is often described as “normalized” if it is in the Third Normal Form because most of the 3NF tables are free of insertion, update, and deletion anomalies.

The relation (Roll_number, Name, Date_of_birth, Age) is

- a) In second normal form but not in third normal form
- b) In third normal form but not in BCNF
- c) In BCNF
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: For the given relation only some of the above FDs are applicable. The applicable FDs are given below:

Date_of_Birth->Age

Name->Roll_number

Roll_number->Name

Finding the closure set of attributes we get the candidate keys:(Roll_number,Date_of_Birth), and (Name,Date_of_Birth) .

On selecting any one of the candidate key we can see that the FD Date_of_Birth->Age is a partial dependency. Hence the relation is in 1NF.

The highest normal form of this relation scheme is

- a) 2NF
- b) 3NF
- c) BCNF
- d) 4NF

[View Answer](#)

Answer: b

Explanation: A super key is a combination of prime attributes and one or more non-prime key attribute(s). It also uniquely identifies a record in a table. Primary key can be defined as super key with minimal attributes.

9. The relation EMPDT1 is defined with attributes empcode(unique), name, street, city, state, and pincode. For any pincode, there is only one city and state. Also, for any given street, city and state, there is just one pincode. In normalization terms EMPDT1 is a relation in

- a) 1NF only
- b) 2NF and hence also in 1NF
- c) 3NF and hence also in 2NF and 1NF
- d) BCNF and hence also in 3NF, 2NF and 1NF

[View Answer](#)

Answer: b

Explanation: Empcode is unique, therefore it is the primary key. Since the primary key consists of a single attribute there will be no partial dependency, hence the relation is in 2NF.

From the question we get the FDs as below:

pincode -> city, state

street,city,state -> pincode

From the FDs we can see that there are transitive dependencies, hence the table is not in 3NF.

10. Which one of the following statements is FALSE?

- a) Any relation with two attributes is in BCNF
- b) A relation in which every key has only one attribute is in 2NF
- c) A prime attribute can be transitively dependent on a key in a 3 NF relation
- d) A prime attribute can be transitively dependent on a key in a BCNF relation

[View Answer](#)

Answer: d

Explanation: A table is in 3NF if and only if, for each of its functional dependencies $X \rightarrow A$, at least one of the following conditions holds:

- * X contains A (that is, $X \rightarrow A$ is trivial functional dependency), or
- * X is a superkey, or
- * A should be prime attribute.

ayzom.com

Database Questions and Answers – Advanced Query Optimization

1. _____ is a procedural extension of Oracle – SQL that offers language constructs similar to those in imperative programming languages.

- a) SQL
- b) PL/SQL
- c) Advanced SQL
- d) PQL

[View Answer](#)

Answer: b

Explanation: PL/SQL is an imperative 3GL that was designed specifically for the seamless processing of SQL commands.

2. _____ combines the data manipulating power of SQL with the data processing power of Procedural languages.

- a) PL/SQL
- b) SQL
- c) Advanced SQL
- d) PQL

[View Answer](#)

Answer: a

Explanation: PL/SQL is an imperative 3GL that was designed specifically for the seamless processing of SQL commands.

3. _____ has made PL/SQL code run faster without requiring any additional work on the part of the programmer.

- a) SQL Server
- b) MySQL
- c) Oracle
- d) SQL Lite

[View Answer](#)

Answer: c

Explanation: An Oracle database is a collection of data treated as a unit. The purpose of a database is to store and retrieve related information.

4. A line of PL/SQL text contains groups of characters known as

- a) Lexical Units
- b) Literals
- c) Textual Units
- d) Identifiers

[View Answer](#)

Answer: a

Explanation: Lexical items can be generally understood to convey a single meaning, much as a lexeme, but are not limited to single words.

5. We use _____ name PL/SQL program objects and units.

- a) Lexical Units
- b) Literals
- c) Delimiters
- d) Identifiers

[View Answer](#)

Answer: d

Explanation: The database object name is referred to as its identifier.

6. A _____ is an explicit numeric, character, string or Boolean value not represented by an identifier.

- a) Comments

- b) Literals
- c) Delimiters
- d) Identifiers

[View Answer](#)

Answer: b

Explanation: The terms literal and constant value are synonymous and refer to a fixed data value.

7. If no header is specified, the block is said to be an _____ PL/SQL block.

- a) Strong
- b) Weak
- c) Empty
- d) Anonymous

[View Answer](#)

Answer: d

Explanation: The terms literal and constant value are synonymous and refer to a fixed data value.

8. _____ is a sequence of zero or more characters enclosed by single quotes.

- a) Integers literal
- b) String literal
- c) String units
- d) String label

[View Answer](#)

Answer: b

Explanation: The terms literal and constant value are synonymous and refer to a fixed data value.

9. In _____ the management of the password for the account can be handled outside of oracle such as operating system.

- a) Database Authentication
- b) Operating System Authentication
- c) Internal Authentication
- d) External Authentication

[View Answer](#)

Answer: b

Explanation: Database management involves the monitoring, administration, and maintenance of the databases and database groups in your enterprise.

10. In _____ of Oracle, the database administrator creates a user account in the database for each user who needs access.

- a) Database Authentication
- b) Operating System Authentication
- c) Internal Authentication
- d) External Authentication

[View Answer](#)

Answer: a

Explanation: Database management involves the monitoring, administration, and maintenance of the databases and database groups in your enterprise.

Database Questions and Answers – Transaction Concept

1. Consider money is transferred from (1)account-A to account-B and (2) account-B to account-A. Which of the following form a transaction?

- a) Only 1
- b) Only 2
- c) Both 1 and 2 individually
- d) Either 1 or 2

[View Answer](#)

Answer: c

Explanation: The term transaction refers to a collection of operations that form a single logical unit of work.

2. A transaction is delimited by statements (or function calls) of the form _____

- a) Begin transaction and end transaction
- b) Start transaction and stop transaction
- c) Get transaction and post transaction
- d) Read transaction and write transaction

[View Answer](#)

Answer: a

Explanation: The transaction consists of all operations executed between the begin transaction and end transaction.

3. Identify the characteristics of transactions

- a) Atomicity
- b) Durability
- c) Isolation
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: Because of the above three properties, transactions are an ideal way of structuring interaction with a database.

4. Which of the following has “all-or-none” property?

- a) Atomicity
- b) Durability
- c) Isolation
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Either all operations of the transaction are reflected properly in the database, or none are.

5. The database system must take special actions to ensure that transactions operate properly without interference from concurrently executing database statements. This property is referred to as

- a) Atomicity
- b) Durability
- c) Isolation
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Even though multiple transactions may execute concurrently, the system guarantees that, for every pair of transactions T_i and T_j , it appears to T_i that either T_j finished execution before T_i started or T_j started execution after T_i finished.

6. The property of a transaction that persists all the crashes is

- a) Atomicity
- b) Durability
- c) Isolation
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: After a transaction completes successfully, the changes it has made to the database persist, even if there are system failures.

7. _____ states that only valid data will be written to the database.

- a) Consistency
- b) Atomicity
- c) Durability
- d) Isolation

[View Answer](#)

Answer: a

Explanation: If for some reason, a transaction is executed that violates the database's consistency rules, the entire transaction will be rolled back and the database will be restored to a state consistent with those rules.

8. Transaction processing is associated with everything below except

- a) Producing detail summary or exception reports
- b) Recording a business activity
- c) Confirming an action or triggering a response
- d) Maintaining a data

[View Answer](#)

Answer: c

Explanation: Collections of operations that form a single logical unit of work are called transactions.

9. The Oracle RDBMS uses the _____ statement to declare a new transaction start and its properties.

- a) BEGIN
- b) SET TRANSACTION
- c) BEGIN TRANSACTION
- d) COMMIT

[View Answer](#)

Answer: b

Explanation: Commit is used to store all the transactions.

10. _____ means that the data used during the execution of a transaction cannot be used by a second transaction until the first one is completed.

- a) Consistency
- b) Atomicity
- c) Durability
- d) Isolation

[View Answer](#)

Answer: d

Explanation: Even though multiple transactions may execute concurrently, the system guarantees that, for every pair of transactions T_i and T_j , it appears to T_i that either T_j finished execution before T_i started or T_j started execution after T_i finished.

Database Questions & Answers – A Simple Transaction Model

1. In SQL, which command is used to issue multiple CREATE TABLE, CREATE VIEW and GRANT statements in a single transaction?

- a) CREATE PACKAGE
- b) CREATE SCHEMA
- c) CREATE CLUSTER
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: A database schema of a database system is its structure described in a formal language supported by the database management system and refers to the organization of data as a blueprint of how a database is constructed.

2. In SQL, the CREATE TABLESPACE is used

- a) To create a place in the database for storage of schema objects, rollback segments, and naming the data files to comprise the tablespace
- b) To create a database trigger
- c) To add/rename data files, to change storage
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Triggers are used to initialize the actions for an activity.

3. Which character function can be used to return a specified portion of a character string?

- a) INSTR
- b) SUBSTRING
- c) SUBSTR
- d) POS

[View Answer](#)

Answer: c

Explanation: SUBSTR are used to match the particular characters in a string.

4. Which of the following is TRUE for the System Variable \$date\$?

- a) Can be assigned to a global variable
- b) Can be assigned to any field only during design time
- c) Can be assigned to any variable or field during run time
- d) Can be assigned to a local variable

[View Answer](#)

Answer: b

Explanation: A database schema of a database system is its structure described in a formal language supported by the database management system and refers to the organization of data as a blueprint of how a database is constructed.

5. What are the different events in Triggers?

- a) Define, Create
- b) Drop, Comment
- c) Insert, Update, Delete
- d) Select, Commit

[View Answer](#)

Answer: c

Explanation: A database trigger is a procedural code that is automatically executed in response to certain events on a particular table or view in a database.

6. Which is the subset of SQL commands used to manipulate Oracle Database Structures, including tables?

- a) Data Definition Language
- b) Data Manipulation Language
- c) Data Described Language
- d) Data Retrieval Language

[View Answer](#)

Answer: a

Explanation: DDL are used to define schema and table characters.

7. The SQL statement SELECT SUBSTR('123456789', INSTR('abcabcabc', 'b'), 4) FROM EMP; prints

- a) 6789
- b) 2345
- c) 1234
- d) 456789

[View Answer](#)

Answer: b

Explanation: SUBSTR are used to match the particular characters in a string.

8. Which of the following SQL command can be used to modify existing data in a database table?

- a) MODIFY
- b) UPDATE
- c) CHANGE
- d) NEW

[View Answer](#)

Answer: b

Explanation: Syntax : UPDATE table_name
SET column1=value1,column2=value2,...
WHERE some_column=some_value; .

9. When SQL statements are embedded inside 3GL, we call such a program as

- a) Nested query
- b) Nested programming
- c) Distinct query
- d) Embedded SQL

[View Answer](#)

Answer: d

Explanation: SQL-99 is the most recent version of standard SQL prescribed by the ANSI.

10. _____ provides option for entering SQL queries as execution time, rather than at the development stage.

- a) PL/SQL
- b) SQL*Plus
- c) SQL
- d) Dynamic SQL

[View Answer](#)

Answer: d

Explanation: Dynamic SQL enables you to write programs that reference SQL statements whose full text is not known until runtime.

Database Questions and Answers – Storage Structure

1. The storage structure which do not survive system crashes are _____

- a) Volatile storage
- b) Non-volatile storage
- c) Stable storage
- d) Dynamic storage

[View Answer](#)

Answer: a

Explanation: Volatile storage, is a computer memory that requires power to maintain the stored information, in other words it needs power to reach the computer memory.

2. Storage devices like tertiary storage, magnetic disk comes under

- a) Volatile storage
- b) Non-volatile storage
- c) Stable storage
- d) Dynamic storage

[View Answer](#)

Answer: b

Explanation: Information residing in nonvolatile storage survives system crashes.

3. For a transaction to be durable, its changes need to be written to _____ storage.

- a) Volatile storage
- b) Non-volatile storage
- c) Stable storage
- d) Dynamic storage

[View Answer](#)

Answer: c

Explanation: Similarly, for a transaction to be atomic, log records need to be written to stable storage before any changes are made to the database on disk.

4. The unit of storage that can store one or more records in a hash file organization are

- a) Buckets
- b) Disk pages
- c) Blocks
- d) Nodes

[View Answer](#)

Answer: a

Explanation: Buckets are used to store one or more records in a hash file organization.

5. A _____ file system is software that enables multiple computers to share file storage while maintaining consistent space allocation and file content.

- a) Storage
- b) Tertiary
- c) Secondary
- d) Cluster

[View Answer](#)

Answer: d

Explanation: With a cluster file system, the failure of a computer in the cluster does not make the file system unavailable.

6. A file produced by a spreadsheet

- a) is generally stored on disk in an ASCII text format
- b) can be used as is by the DBMS
- c) all of the mentioned
- d) none of the mentioned

[View Answer](#)

Answer: a

Explanation: ASCII text format uses the standard text file for the changing the value.

7. SQL means _____

- a) Storage Discrete Language
- b) Storage Definition Language
- c) Storage Definition Localisation
- d) Storage Discrete Localisation

[View Answer](#)

Answer: b

Explanation: It specifies internal schema and also mapping between two schemas.

8. Which of the following are the process of selecting the data storage and data access characteristics of the database?

- a) Logical database design
- b) Physical database design
- c) Testing and performance tuning
- d) Evaluation and selecting

[View Answer](#)

Answer: b

Explanation: Physical database design is the process of selecting the data storage and data access characteristics of the database.

9. Which of the following is the oldest database model?

- a) Relational
- b) Hierarchical
- c) Physical
- d) Network

[View Answer](#)

Answer: d

Explanation: Network model has data stored in a hierarchical network flow.

10. The process of saving information onto secondary storage devices is referred to as

- a) Backing up
- b) Restoring
- c) Writing
- d) Reading

[View Answer](#)

Answer: c

Explanation: The information is written into the secondary storage device.

Database Questions and Answers – Transaction Atomicity and Durability

1. A transaction may not always complete its execution successfully. Such a transaction is termed

- a) Aborted
- b) Terminated
- c) Closed
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: If we are to ensure the atomicity property, an aborted transaction must have no effect on the state of the database.

2. If an transaction is performed in a database and committed, the changes are taken to the previous state of transaction by

- a) Flashback
- b) Rollback
- c) Both Flashback and Rollback
- d) Cannot be done

[View Answer](#)

Answer: d

Explanation: Once committed the changes cannot be rolled back.

3. Each modification done in database transaction are first recorded into the

- a) Harddrive
- b) Log
- c) Disk
- d) Datamart

[View Answer](#)

Answer: b

Explanation: After commit is issued the data are stored in a database and stored in drive.

4. When the transaction finishes the final statement the transaction enters into

- a) Active state
- b) Committed state
- c) Partially committed state
- d) Abort state

[View Answer](#)

Answer: c

Explanation: The commit statement has to be issued to enter into committed state.

5. The name of the transaction file shall be provided by the operator and the file that contains the edited transactions ready for execution shall be called

- a) Batch. Exe
- b) Trans. Exe
- c) Opt. Exe
- d) Edit.Exe

[View Answer](#)

Answer: c

Explanation: Transactions has to be managed by the executable files.

6. Which of the following is an atomic sequence of database actions?

- a) Transaction

- b) Concurrency
- c) Relations
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Transaction is a collection of operations that provides single logical function in database.

7. If the state of the database no longer reflects a real state of the world that the database is supposed to capture, then such a state is called

- a) Consistent state
- b) Parallel state
- c) Atomic state
- d) Inconsistent state

[View Answer](#)

Answer: d

Explanation: If the state of the database no longer reflects a real state of the world that the database is supposed to capture, then such a state is called a consistent state.

8. _____ means that data used during the execution of a transaction cannot be used by a second transaction until the first one is completed.

- a) Serializability
- b) Atomicity
- c) Isolation
- d) Time stamping

[View Answer](#)

Answer: c

Explanation: Isolation means that data used during the execution of a transaction can't be used by a second transaction until the first one is completed.

9. DBMS periodically suspends all processing and synchronizes its files and journals through the use of

- a) Checkpoint facility
- b) Backup facility
- c) Recovery manager
- d) Database change log

[View Answer](#)

Answer: a

Explanation: DBMS periodically suspends all processing and synchronizes its files and journals though the use of Check point facility.

10. Which of the following is not a state in transaction?

- a) Active
- b) Terminated
- c) Aborted
- d) Partially committed

[View Answer](#)

Answer: b

Explanation: The transaction states are abort, active, committed, partially committed, Failed.

Database Questions and Answers – Querying database part 3

1. _____ joins are SQL server default

- a) Outer
- b) Inner
- c) Equi
- d) None of the Mentioned

[View Answer](#)

Answer: b

Explanation: Inner query joins only the rows that are matching.

2. The _____ is essentially used to search for patterns in target string.

- a) Like Predicate
- b) Null Predicate
- c) In Predicate
- d) Out Predicate

[View Answer](#)

Answer: a

Explanation: Like matches the pattern with the query.

3. Which of the following is/are the Database server functions?

- i) Data management
- ii) Transaction management
- iii) Compile queries
- iv) Query optimization

- a) i, ii, and iv only
- b) i, ii and iii only
- c) ii, iii and iv only
- d) All i, ii, iii, and iv

[View Answer](#)

Answer: a

Explanation: All these are functions of the database.

4. To delete a database _____ command is used

- a) Delete database database_name
- b) Delete database_name
- c) drop database database_name
- d) drop database_name

[View Answer](#)

Answer: c

Explanation: This will delete the database with its structure.

5. _____ is a combination of two or more attributes used as a primary key

- a) Composite Key
- b) Alternate Key
- c) Candidate Key
- d) Foreign Key

[View Answer](#)

Answer: a

Explanation: Primary keys together form the composite key.

6. Which of the following is not the function of client?

- a) Compile queries
- b) Query optimization
- c) Receive queries
- d) Result formatting and presentation

[View Answer](#)

Answer: b

Explanation: Query optimization is used to improve quality.

7. _____ is a special type of stored procedure that is automatically invoked whenever the data in the table is modified.

- a) Procedure
- b) Trigger
- c) Cursor
- d) None of the Mentioned

[View Answer](#)

Answer: b

Explanation: Triggers are used to initiate an action to take place.

8. _____ requires that data should be made available to only authorized users.

- a) Data integrity
- b) Privacy
- c) Security
- d) None of the Mentioned

[View Answer](#)

Answer: c

Explanation: Some algorithms may be used for the security.

9. Some of the utilities of DBMS are _____

- i) Loading ii) Backup iii) File organization iv) Process Organization
- a) i, ii, and iv only
- b) i, ii and iii only
- c) ii, iii and iv only
- d) All i, ii, iii, and iv

[View Answer](#)

Answer: b

Explanation: Processing is not the a utility in dbms.

10. _____ allows individual row operation to be performed on a given result set or on the generated by a selected by a selected statement.

- a) Procedure
- b) Trigger
- c) Cursor
- d) None of the Mentioned

[View Answer](#)

Answer: c

Explanation: Triggers are used to initiate an action to take place.

Database Questions and Answers – Querying Database – 4

1. Which is essential a business problem not a data problem:

- a) Data
- b) Database
- c) Database design
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: SQL-99 is the most recent version of standard SQL prescribed by the ANSI.

2. Which is primarily the result of a thorough understanding of information about an enterprise:

- a) Data
- b) Database
- c) Database design
- d) Data modeling

[View Answer](#)

Answer: d

Explanation: Data modelling designs the data in a secured manner.

3. McFadden has defined normalization in his book _____

- a) Database modern management
- b) Management database of modern
- c) Modern database management
- d) Database management

[View Answer](#)

Answer: c

Explanation: SQL-99 is the most recent version of standard SQL prescribed by the ANSI.

4. The database design prevents some data from being represented due to _____

- a) Deletion anomalies
- b) Insertion anomalies
- c) Update anomaly
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Insertion anomaly is due to confusion in data deletion or insertion.

5. How many types of insertion anomalies:

- a) 1
- b) 2
- c) 3
- d) 4

[View Answer](#)

Answer: b

Explanation: Insertion anomaly is due to confusion in data deletion or insertion.

6. Who developed the normalization process:

- a) E.F. codd
- b) F.F. codd

- c) E.E. codd
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Normalization helps in improving the quality of the data.

7. E.F.Codd developed the normalization process in the which early:

- a) 1969
- b) 1970
- c) 1971
- d) 1972

[View Answer](#)

Answer: b

Explanation: Normalization helps in improving the quality of the data.

8. Which is a bottom-up approach to database design that design by examining the relationship between attributes:

- a) Functional dependency
- b) Database modeling
- c) Normalization
- d) Decomposition

[View Answer](#)

Answer: c

Explanation: Normalization helps in improving the quality of the data.

9. Which is the process of breaking a relation into multiple relations:

- a) Functional dependency
- b) Database modeling
- c) Normalization
- d) Decomposition

[View Answer](#)

Answer: d

Explanation: SQL-99 is the most recent version of standard SQL prescribed by the ANSI.

10. Which formal method that locates and analyses relation schemas on the basis of their primary, candidate keys, and the FD's that are present among the attributes of these schemas:

- a) Functional dependency
- b) Database modeling
- c) Normalization
- d) Decomposition

[View Answer](#)

Answer: c

Explanation: Normalization helps in improving the quality of the data.

Database Questions and Answers – Querying database part 5

1. Which is refers to a stalemate situation due to which no further progress is possible as computer await response of each other:

- a) Concurrency
- b) Deadlock
- c) Backup
- d) Recovery

[View Answer](#)

Answer: b

Explanation: Deadlock will stop further processing.

2. Which is a duplicate copy of a file program that is stored on a different storage media than the original location:

- a) Concurrency
- b) Deadlock
- c) Backup
- d) Recovery

[View Answer](#)

Answer: c

Explanation: Backup is required to take all the data.

3. Which is the duplication of computer operations and routine backups to combat any unforeseen problems:

- a) Concurrency
- b) Deadlock
- c) Backup
- d) Recovery

[View Answer](#)

Answer: d

Explanation: Recovery means to take the backup data while there is a crash.

4. Optimization that is basically related to the rewriter module is termed as _____

- a) Semantic query optimization
- b) Global query optimization
- c) All of the Mentioned
- d) None of the Mentioned

[View Answer](#)

Answer: a

Explanation: SQL-99 is the most recent version of standard SQL prescribed by the ANSI.

5. Optimization basically related to the Rewrite module is termed as _____

- a) Semantic query optimization
- b) Global query optimization
- c) All of the Mentioned
- d) None of the Mentioned

[View Answer](#)

Answer: a

Explanation: SQL-99 is the most recent version of standard SQL prescribed by the ANSI.

6. Database security helps organizations to protect data from _____

- a) Internal users
- b) External users

- c) Non-external users
- d) Non internal users

[View Answer](#)

Answer: b

Explanation: External users are the people who do not involve in the processing of the database.

7. Copying files to secondary or specific devices is known as _____

- a) Retrieve
- b) Backup
- c) Recovery
- d) Deadlock

[View Answer](#)

Answer: b

Explanation: Backup is required to take all the data.

8. How many types of recovery control techniques:

- a) 2
- b) 3
- c) 4
- d) 5

[View Answer](#)

Answer: a

Explanation: Recovery means to take the backup data while there is a crash.

9. Which are types of recovery control techniques:

- a) Deferred update
- b) Immediate update
- c) All of the Mentioned
- d) None of the Mentioned

[View Answer](#)

Answer: c

Explanation: Recovery means to take the backup data while there is a crash.

10. Which server can join the indexes when only multiple indexes combined can cover the query:

- a) SQL
- b) DBMS
- c) RDBMS
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Indexing reduces the difficulty in searching the data.

Database Questions and Answers – Implementation of Isolation Levels

1. In concurrency control policy the lock is obtained on

- a) Entire database
- b) A particular transaction alone
- c) All the new elements
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: It is to avoid deadlock.

2. A concurrency-control policy such as this one leads to _____ performance since it forces transactions to wait for preceding transactions to finish before they can start.

- a) Good
- b) Average
- c) Poor
- d) Unstable

[View Answer](#)

Answer: c

Explanation: It provides a poor degree of concurrency.

3. _____ are used to ensure that transactions access each data item in order of the transactions' _____ if their accesses conflict.

- a) Zone
- b) Relay
- c) Line
- d) Timestamps

[View Answer](#)

Answer: d

Explanation: When this is not possible, offending transactions are aborted and restarted with a new timestamp.

[View Answer](#)

Answer: c

Explanation: View is the temporary space created for the database.

5. In which scenario would you use the ROLLUP operator for expression or columns within a GROUP BY clause?

- a) To find the groups forming the subtotal in a row
- b) To create group-wise grand totals for the groups specified within a GROUP BY clause
- c) To create a grouping for expressions or columns specified within a GROUP BY clause in one direction, from right to left for calculating the subtotals
- d) To create a grouping for expressions or columns specified within a GROUP BY clause in all possible directions, which is cross-tabular report for calculating the subtotals

[View Answer](#)

Answer: c

Explanation: View is the temporary space created for the database.

Which rows would be made permanent in the CUST table?

- a) row 4 only
- b) rows 2 and 4
- c) rows 3 and 4
- d) rows 1 and 4

[View Answer](#)

Answer: c

Explanation: View is the temporary space created for the database.

7. Which statement is true regarding external tables?

- a) The default REJECT LIMIT for external tables is UNLIMITED
- b) The data and metadata for an external table are stored outside the database
- c) ORACLE_LOADER and ORACLE_DATAPUMP have exactly the same functionality when used with an external table
- d) The CREATE TABLE AS SELECT statement can be used to unload data into regular table in the database from an external table

[View Answer](#)

Answer: d

Explanation: This will replicate the table as in the select statement.

8. A non-correlated subquery can be defined as _____

- a) A set of sequential queries, all of which must always return a single value
- b) A set of sequential queries, all of which must return values from the same table
- c) A SELECT statement that can be embedded in a clause of another SELECT statement only
- d) A set of one or more sequential queries in which generally the result of the inner query is used as the search value in the outer query

[View Answer](#)

Answer: d

Explanation: This will replicate the table as in the select statement.

Which statement is true regarding the above FLASHBACK operation?

- a) It recovers only the first DEPT table
- b) It recovers only the second DEPT table
- c) It does not recover any of the tables because FLASHBACK is not possible in this case
- d) It recovers both the tables but the names would be changed to the ones assigned in the RECYCLEBIN

[View Answer](#)

Answer: b

Explanation: This will replicate the table as in the select statement.

What would be the outcome of the above query?

- a) It would not display any values
- b) It would display the value TWO once
- c) It would display the value TWO twice
- d) It would display the values ONE, TWO, and TWO

[View Answer](#)

Answer: c

Explanation: This will replicate the table as in the select statement.

Database Questions and Answers – Transactions as SQL Statements

1. Which of the following is not a property of transactions?

- a) Atomicity
- b) Concurrency
- c) Isolation
- d) Durability

[View Answer](#)

Answer: d

Explanation: ACID properties are the properties of transactions.

2. SNAPSHOT is used for (DBA)

- a) Synonym
- b) Tablespace
- c) System server
- d) Dynamic data replication

[View Answer](#)

Answer: d

Explanation: Snapshot gets the instance of the database at that time.

3. Isolation of the transactions is ensured by

- a) Transaction management
- b) Application programmer
- c) Concurrency control
- d) Recovery management

[View Answer](#)

Answer: c

Explanation: ACID properties are the properties of transactions.

4. Constraint checking can be disabled in existing _____ and _____ constraints so that any data you modify or add to the table is not checked against the constraint.

- a) CHECK, FOREIGN KEY
- b) DELETE, FOREIGN KEY
- c) CHECK, PRIMARY KEY
- d) PRIMARY KEY, FOREIGN KEY

[View Answer](#)

Answer: a

Explanation: Check and foreign constraints are used to constraint the table data.

5. Problems occurs if we don't implement a proper locking strategy

- a) Dirty reads
- b) Phantom reads
- c) Lost updates
- d) Unrepeatable reads

[View Answer](#)

Answer: d

Explanation: In a concurrent execution of these transactions, it is intuitively clear that they conflict, but this is a conflict not captured by our simple model. This situation is referred to as the phantom phenomenon, because a conflict may exist on "phantom" data.

6. Which of the following fixed database roles can add or remove user IDs?

- a) db_accessadmin
- b) db_securityadmin
- c) db_setupadmin
- d) db_sysadmin

[View Answer](#)

Answer: a

Explanation: The database can be accessed by assigning the roles.

7. By default sql server has _____ isolation level

- a) READ COMMITTED
- b) READ UNCOMMITTED
- c) SERIALIZABLE
- d) REPEATABLE READ

[View Answer](#)

Answer: a

Explanation: Read committed is used to commit the default read operation.

8. Which of the following statements is/are not true for SQL profiler?

- a) Enables you to monitor events
- b) Check if rows are being inserted properly
- c) Check the performance of a stored procedure
- d) ALL of the mentioned

[View Answer](#)

Answer: c

Explanation: Read committed is used to commit the default read operation.

9. Which of the following is the original purpose of SQL?

- a) To specify the syntax and semantics of SQL data definition language
- b) To specify the syntax and semantics of SQL manipulation language
- c) To define the data structures
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: Read committed is used to commit the default read operation.

10. SQL can be used to:

- a) Create database structures only
- b) Query database data only
- c) Modify database data only
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: In a concurrent execution of these transactions, it is intuitively clear that they conflict, but this is a conflict not captured by our simple model. This situation is referred to as the phantom phenomenon, because a conflict may exist on “phantom” data.

Database Questions and Answers – Lock-Based Protocols

1. In order to maintain transactional integrity and database consistency, what technology does a DBMS deploy?

- a) Triggers
- b) Pointers
- c) Locks
- d) Cursors

[View Answer](#)

Answer: c

Explanation: Locks are used to maintain database consistency.

2. A lock that allows concurrent transactions to access different rows of the same table is known as a

- a) Database-level lock
- b) Table-level lock
- c) Page-level lock
- d) Row-level lock

[View Answer](#)

Answer: d

Explanation: Locks are used to maintain database consistency.

3. Which of the following are introduced to reduce the overheads caused by the log-based recovery?

- a) Checkpoints
- b) Indices
- c) Deadlocks
- d) Locks

[View Answer](#)

Answer: a

Explanation: Checkpoints are introduced to reduce overheads caused by the log-based recovery.

4. Which of the following protocols ensures conflict serializability and safety from deadlocks?

- a) Two-phase locking protocol
- b) Time-stamp ordering protocol
- c) Graph based protocol
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Time-stamp ordering protocol ensures conflict serializability and safety from deadlocks.

5. Which of the following is the block that is not permitted to be written back to the disk?

- a) Dead code
- b) Read only
- c) Pinned
- d) Zapped

[View Answer](#)

Answer: c

Explanation: A block that is not permitted to be written back to the disk is called pinned.

6. If transaction T_i gets an explicit lock on the file F_c in exclusive mode, then it has an _____ on all the records belonging to that file.

- a) Explicit lock in exclusive mode
- b) Implicit lock in shared mode

- c) Explicit lock in shared mode
- d) Implicit lock in exclusive mode

[View Answer](#)

Answer: d

Explanation: If transaction T_i gets an explicit lock on the file F_c in exclusive mode, then it has an implicit lock in exclusive mode on all the records belonging to that file.

7. Which refers to a property of computer to run several operation simultaneously and possible as computers await response of each other

- a) Concurrency

b) Deadlock

c) Backup

d) Recovery

[View Answer](#)

Answer: a

Explanation: Concurrency is a property of systems in which several computations are executing simultaneously, and potentially interacting with each other.

8. All lock information is managed by a _____ which is responsible for assigning and policing the locks used by the transactions.

a) Scheduler

b) DBMS

c) Lock manager

d) Locking agent

[View Answer](#)

Answer: c

Explanation: A distributed lock manager (DLM) provides distributed software applications with a means to synchronize their accesses to shared resources.

9. The ____ lock allows concurrent transactions to access the same row as long as they require the use of different fields within that row.

a) Table-level

b) Page-level

c) Row-level

d) Field-level

[View Answer](#)

Answer: d

Explanation: Lock is limited to the attributes of the relation.

10. Which of the following is a procedure for acquiring the necessary locks for a transaction where all necessary locks are acquired before any are released?

a) Record controller

b) Exclusive lock

c) Authorization rule

d) Two phase lock

[View Answer](#)

Answer: d

Explanation: Two-phase lock is a procedure for acquiring the necessary locks for a transaction where all necessary locks are acquired before any are released.

Database Questions and Answers – Deadlocks

1. A system is in a _____ state if there exists a set of transactions such that every transaction in the set is waiting for another transaction in the set.

- a) Idle
- b) Waiting
- c) Deadlock
- d) Ready

[View Answer](#)

Answer: c

Explanation: When one data item is waiting for another data item in a transaction then system is in deadlock.

2. The deadlock state can be changed back to stable state by using _____ statement.

- a) Commit
- b) Rollback
- c) Savepoint
- d) Deadlock

[View Answer](#)

Answer: b

Explanation: Rollback is used to rollback to the point before lock is obtained.

3. What are the ways of dealing with deadlock?

- a) Deadlock prevention
- b) Deadlock recovery
- c) Deadlock detection
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: Deadlock prevention is also called as deadlock recovery. Prevention is commonly used if the probability that the system would enter a deadlock state is relatively high; otherwise, detection and recovery are more efficient.

4. When transaction T_i requests a data item currently held by T_j , T_i is allowed to wait only if it has a timestamp smaller than that of T_j (that is, T_i is older than T_j). Otherwise, T_i is rolled back (dies). This is

- a) Wait-die
- b) Wait-wound
- c) Wound-wait
- d) Wait

[View Answer](#)

Answer: a

Explanation: The wait-die scheme is a non-preemptive technique.

5. When transaction T_i requests a data item currently held by T_j , T_i is allowed to wait only if it has a timestamp larger than that of T_j (that is, T_i is younger than T_j). Otherwise, T_j is rolled back (T_j is wounded by T_i). This is

- a) Wait-die
- b) Wait-wound
- c) Wound-wait
- d) Wait

[View Answer](#)

Answer: c

Explanation: The wound-wait scheme is a preemptive technique. It is a counterpart to the wait-die scheme.

6. The situation where the lock waits only for a specified amount of time for another lock to be released is

- a) Lock timeout
- b) Wait-wound
- c) Timeout
- d) Wait

[View Answer](#)

Answer: a

Explanation: The timeout scheme is particularly easy to implement, and works well if transactions are short and if longwaits are likely to be due to deadlocks.

7. The deadlock in a set of a transaction can be determined by

- a) Read-only graph
- b) Wait graph
- c) Wait-for graph
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Each transaction involved in the cycle is said to be deadlocked.

8. A deadlock exists in the system if and only if the wait-for graph contains a _____

- a) Cycle
- b) Direction
- c) Bi-direction
- d) Rotation

[View Answer](#)

Answer: a

Explanation: Each transaction involved in the cycle is said to be deadlocked.

9. Selecting the victim to be rollbacked to the previous state is determined by the minimum cost. The factors determining cost of rollback is

- a) How long the transaction has computed, and how much longer the transaction will compute before it completes its designated task
- b) How many data items the transaction has used
- c) How many more data items the transaction needs for it to complete
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: We should roll back those transactions that will incur the minimum cost.

10. _____ rollback requires the system to maintain additional information about the state of all the running transactions.

- a) Total
- b) Partial
- c) Time
- d) Commit

[View Answer](#)

Answer: b

Explanation: In total rollback abort the transaction and then restart it.

Database Questions and Answers – Multiple Granularity

1. In a granularity hierarchy the highest level represents the

- a) Entire database
- b) Area
- c) File
- d) Record

[View Answer](#)

Answer: a

Explanation: This level is the root of the tree.

2. In a database the file is contained in _____

- a) Entire database
- b) Two area
- c) One area
- d) more than one area

[View Answer](#)

Answer: c

Explanation: This level is below the root of the tree.

3. If a node is locked in an intention mode, explicit locking is done at a lower level of the tree. This is called

- a) Intention lock modes
- b) Explicit lock
- c) Implicit lock
- d) Exclusive lock

[View Answer](#)

Answer: a

Explanation: There is an intention mode associated with shared mode, and there is one with an exclusive mode.

4. If a node is locked in _____ explicit locking is being done at a lower level of the tree, but with only shared-mode locks.

- a) Intention lock modes
- b) Intention-shared-exclusive mode
- c) Intention-exclusive (IX) mode
- d) Intention-shared (IS) mode

[View Answer](#)

Answer: a

Explanation: There is an intention mode associated with shared mode, and there is one with an exclusive mode.

5. If a node is locked in _____ then explicit locking is being done at a lower level, with exclusive-mode or shared-mode locks.

- a) Intention lock modes
- b) Intention-shared-exclusive mode
- c) Intention-exclusive (IX) mode
- d) Intention-shared (IS) mode

[View Answer](#)

Answer: c

Explanation: There is an intention mode associated with shared mode, and there is one with an exclusive mode.

6. If a node is locked in _____ the subtree rooted by that node is locked explicitly in shared mode, and that explicit locking is being done at a lower level with exclusive-mode locks.

- a) Intention lock modes

- b) shared and intention-exclusive (SIX) mode
- c) Intention-exclusive (IX) mode
- d) Intention-shared (IS) mode

[View Answer](#)

Answer: b

Explanation: There is an intention mode associated with shared mode, and there is one with an exclusive mode.

7. _____ denotes the largest timestamp of any transaction that executed write(Q) successfully.

- a) W-timestamp(Q)
- b) R-timestamp(Q)
- c) RW-timestamp(Q)
- d) WR-timestamp(Q)

[View Answer](#)

Answer: a

Explanation: The most common method for doing ordering transaction is to use a timestamp-ordering scheme.

8. The _____ ensures that any conflicting read and write operations are executed in timestamp order.

- a) Timestamp-ordering protocol
- b) Timestamp protocol
- c) W-timestamp
- d) R-timestamp

[View Answer](#)

Answer: a

Explanation: The most common method for doing ordering transaction is to use a timestamp-ordering scheme.

9. The _____ requires that each transaction T_i executes in two or three different phases in its lifetime, depending on whether it is a read-only or an update transaction.

- a) Validation protocol
- b) Validation-based protocol
- c) Timestamp protocol
- d) Timestamp-ordering protocol

[View Answer](#)

Answer: a

Explanation: A concurrency-control scheme imposes the overhead of code execution and possible delay of transactions. It may be better to use an alternative scheme that imposes less overhead.

10. This validation scheme is called the _____ scheme since transactions execute optimistically, assuming they will be able to finish execution and validate at the end.

- a) Validation protocol
- b) Validation-based protocol
- c) Timestamp protocol
- d) Optimistic concurrency-control

[View Answer](#)

Answer: a

Explanation: A concurrency-control scheme imposes the overhead of code execution and possible delay of transactions. It may be better to use an alternative scheme that imposes less overhead.

Database Questions and Answers – Multiversion Schemes

1. The most recent version of standard SQL prescribed by the American National Standards Institute is

- a) SQL 2016
- b) SQL 2002
- c) SQL – 4
- d) SQL2

[View Answer](#)

Answer: a

Explanation: SQL-2016 is the most recent version of standard SQL prescribed by the ANSI.

2. ANSI-standard SQL allows the use of special operators in conjunction with the WHERE clause. A special operator used to check whether an attribute value is null is

- a) BETWEEN
- b) IS NULL
- c) LIKE
- d) IN

[View Answer](#)

Answer: b

Explanation: Exists is used to check whether an attribute value is null or not in conjunction with the where clause.

3. A lock that prevents the use of any tables in the database from one transaction while another transaction is being processed is called a

- a) Database-level lock
- b) Table-level lock
- c) Page-level lock
- d) Row-level lock

[View Answer](#)

Answer: a

Explanation: Data base-level lock prevents the use of any tables in the data base from one transaction while another transaction is being processed.

4. A condition that occurs when two transactions wait for each other to unlock data is known as a(n)

- a) Shared lock
- b) Exclusive lock
- c) Binary lock
- d) Deadlock

[View Answer](#)

Answer: d

Explanation: Deadlock occurs when two transactions wait for each other to unlock data.

5. _____ means that data used during the execution of a transaction cannot be used by a second transaction until the first one is completed.

- a) Serializability
- b) Atomicity
- c) Isolation
- d) Time stamping

[View Answer](#)

Answer: c

Explanation: Isolation means that data used during the execution of a transaction can't be used by a second transaction until the first one is completed.

6. A unit of storage that can store one or more records in a hash file organization is denoted as

- a) Buckets
- b) Disk pages
- c) Blocks
- d) Nodes

[View Answer](#)

Answer: a

Explanation: Buckets are used to store one or more records in a hash file organization.

7. The file organization which allows us to read records that would satisfy the join condition by using one block read is

- a) Heap file organization
- b) Sequential file organization
- c) Clustering file organization
- d) Hash files organization

[View Answer](#)

Answer: c

Explanation: Clustering file organization allows us to read records that would satisfy the join condition by using one block read.

8. Which of the following is not true about B+ trees?

- a) B+ tree index takes the form of balanced tree
- b) Performance of B+ tree degrades as the file grows
- c) Look-up in B+ tree is straightforward and efficient
- d) Insertion and deletion in B+ tree is complicated but efficient

[View Answer](#)

Answer: b

Explanation: The answer is evident.

9. The extent of the database resource that is included with each lock is called the level of

- a) Impact
- b) Granularity
- c) Management
- d) DBMS control

[View Answer](#)

Answer: b

Explanation: The extent of the data base resource that is included with each lock is called the level of Granularity.

10. DBMS periodically suspends all processing and synchronizes its files and journals through the use of

- a) Checkpoint facility
- b) Backup facility
- c) Recovery manager
- d) Database change log

[View Answer](#)

Answer: a

Explanation: DBMS periodically suspends all processing and synchronizes its files and journals through the use of Check point facility.

Database Questions and Answers – Snapshot Isolation

1. Snapshot isolation is a particular type of _____ scheme.

- a) Concurrency-control
- b) Concurrency-allowance
- c) Redirection
- d) Repetition-allowance

[View Answer](#)

Answer: a

Explanation: It has gained wide acceptance in commercial and open-source systems, including Oracle, PostgreSQL, and SQL Server.

2. Snapshot isolation is used to give

- a) Transaction a snapshot of the database
- b) Database a snapshot of the transaction
- c) Database a snapshot of committed values in the transaction
- d) Transaction a snapshot of the database and Database a snapshot of committed values in the transaction

[View Answer](#)

Answer: d

Explanation: The data values in the snapshot consist only of values written by committed transactions.

3. Lost update problem is

- a) Second update overwrites the first
- b) First update overwrites the second
- c) The updates are lost due to conflicting problem
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Lost update problem has to be resolved.

4. Under first updaters wins the system uses a _____ mechanism that applies only to updates.

- a) Close
- b) Read
- c) Locking
- d) Beat

[View Answer](#)

Answer: c

Explanation: Reads are unaffected by this, since they do not obtain locks.

5. When a transaction T_i attempts to update a data item, it requests a _____ on that data item.

- a) Read lock
- b) Update lock
- c) Write lock
- d) Chain lock

[View Answer](#)

Answer: c

Explanation: Reads are unaffected by this, since they do not obtain locks.

6. Each of a pair of transactions has read data that is written by the other, but there is no data written by both transactions, is referred to as

- a) Read skew
- b) Update skew

- c) Write lock
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Write skew is the issue addressed here.

7. An application developer can guard against certain snapshot anomalies by appending a _____ clause to the SQL select query.

- a) For update
- b) For read
- c) For write
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Adding the for update clause causes the system to treat data that are read as if they had been updated for purposes of concurrency control.

Which statement is true regarding the PROD_ID_PK constraint?

- a) It would be created only if a unique index is manually created first
- b) It would be created and would use an automatically created unique index
- c) It would be created and would use an automatically created no unique index
- d) It would be created and remains in a disabled state because no index is specified in the command

[View Answer](#)

Answer: b

Explanation: Syntax: create table table_name(name constraint).

The sequence SEQ1 has generated numbers up to the maximum limit of 200. You issue the following SQL statement:

SELECT seq1.nextval FROM dual;

What is displayed by the SELECT statement?

- a) 1
- b) 10
- c) 100
- d) an error

[View Answer](#)

Answer: a

Explanation: Sequence is used to generate a series of values.

10. In which scenario would you use the ROLLUP operator for expression or columns within a GROUP BY clause?

- a) To find the groups forming the subtotal in a row
- b) To create group-wise grand totals for the groups specified within a GROUP BY clause
- c) To create a grouping for expressions or columns specified within a GROUP BY clause in one direction, from right to left for calculating the subtotals
- d) To create a grouping for expressions or columns specified within a GROUP BY clause in all possible directions, which is cross-tabular report for calculating the subtotals

[View Answer](#)

Answer: c

Explanation: Sequence is used to generate a series of values.

Database Questions and Answers – Insertion Deletion Predicate Reads

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

1. Which statements are correct regarding indexes?

- a) When a table is dropped, the corresponding indexes are automatically dropped
- b) For each DML operation performed, the corresponding indexes are automatically updated
- c) A non-deferrable PRIMARY KEY or UNIQUE KEY constraint in a table automatically creates a unique index
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: Indexes are used to access the data efficiently.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

The DELETE statement results in the following error:

ERROR at line 1: table or view does not exist

What would be the outcome?

- a) All the statements before the DELETE statement would be rolled back
- b) All the statements before the DELETE statement would be implicitly committed within the session
- c) All the statements up to the ALTER TABLE statement would be committed and the outcome of UPDATE statement would be rolled back
- d) All the statements up to the ALTER TABLE statement would be committed and the outcome of the UPDATE statement is retained uncommitted within the session

[View Answer](#)

Answer: d

Explanation: Committing a transaction refers to making the changes to record in the database.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
```

VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;

What would be the outcome of the above query?

- a) It would not display any values
- b) It would display the value TWO once
- c) It would display the value TWO twice
- d) It would display the values ONE, TWO, and TWO

[View Answer](#)

Answer: c

Explanation: The VERSIONS BETWEEN clause of the SELECT statement is used to create a Flashback Version Query.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

4. A non-correlated subquery can be defined as _____

- a) A set of sequential queries, all of which must always return a single value
- b) A set of sequential queries, all of which must return values from the same table
- c) A SELECT statement that can be embedded in a clause of another SELECT statement only
- d) A set of one or more sequential queries in which generally the result of the inner query is used as the search value in the outer query

[View Answer](#)

Answer: d

Explanation: A noncorrelated subquery is subquery that is independent of the outer query and it can be executed on its own without relying on main outer query.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

5. Which statement is true regarding synonyms?

- a) Synonyms can be created for tables but not views
- b) Synonyms are used to reference only those tables that are owned by another user
- c) A public synonym and a private synonym can exist with the same name for the same table
- d) The DROP SYNONYM statement removes the synonym, and the status of the table on which the synonym has been created becomes invalid

[View Answer](#)

Answer: c

Explanation: A synonym is an alias or alternate name for a table, view, sequence, or other schema object. They are used mainly to make it easy for users to access database objects owned by other users.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

Which statement is true regarding the execution of the above commands?

- a) Statement 1 would not execute because the WITH GRANT option is missing
- b) Statement 1 would not execute because the IDENTIFIED BY clause is missing
- c) Statement 3 would not execute because role and system privileges cannot be granted together in a single GRANT statement
- d) Statement 2 would not execute because system privileges and object privileges cannot be granted together in a single GRANT command

[View Answer](#)

Answer: d

Explanation: The GRANT statement is used to give privileges to a specific user or role, or to all users, to perform actions on database objects.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

What would be the outcome after executing the statements?

- a) SCOTT would be able to query the OE.ORDERS table
- b) SCOTT would not be able to query the OE.ORDERS table
- c) The REVOKE statement would remove the SELECT privilege from SCOTT as well as from the role R1
- d) The REVOKE statement would give an error because the SELECT privilege has been granted to the role R1

[View Answer](#)

Answer: a

Explanation: The REVOKE statement is used to remove privileges from a specific user or role, or from all users, to perform actions on database objects.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

Which statement describes the consequences?

- a) No SQL statement would be rolled back
- b) Both the DELETE statements would be rolled back
- c) Only the second DELETE statement would be rolled back
- d) Both the DELETE statements and the UPDATE statement would be rolled back

[View Answer](#)

Answer: d

Explanation: The SAVEPOINT statement names and marks the current point in the processing of a transaction. With the ROLLBACK TO statement, savepoints undo parts of a transaction instead of the whole transaction.

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

Which statements are true regarding the above view?

- a) It allows you to insert details of all new staff into the EMPLOYEES table
- b) It allows you to delete the details of the existing sales staff from the EMPLOYEES table
- c) It allows you to update the job ids of the existing sales staff to any other job id in the EMPLOYEES table
- d) It allows you to insert the IDs, last

[View Answer](#)

Answer: d

Explanation: SQL Create view syntax :

```
CREATE VIEW view_name AS
SELECT column_name(s)
FROM TABLE_NAME
WHERE condition.
```

3. Evaluate the following statements:

```
CREATE TABLE digits
(id NUMBER(2),
description VARCHAR2(15));
INSERT INTO digits VALUES (1,'ONE');
UPDATE digits SET description ='TWO'WHERE id=1;
INSERT INTO digits VALUES (2 , 'TWO');
COMMIT;
DELETE FROM digits;
SELECT description FROM digits
VERSIONS BETWEEN TIMESTAMP MINVALUE AND MAXVALUE;
```

[View Answer](#)

Answer: c

Explanation: External tables are created using the SQL CREATE TABLE...ORGANIZATION EXTERNAL statement. When an external table is created, you specify type ,default directory, access parameters and location.

Database Questions and Answers – Concurrency in Index Structures

1. The method of access that uses key transformation is called as

- a) Direct
- b) Hash
- c) Random
- d) Sequential

[View Answer](#)

Answer: b

Explanation: Hash technique uses particular hash key value.

2. Why do we need concurrency control on B+ trees ?

- a) To remove the unwanted data
- b) To easily add the index elements
- c) To maintain accuracy of index
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Indices do not have to be treated like other database structures.

3. How many techniques are available to control concurrency on B+ trees?

- a) One
- b) Three
- c) Four
- d) None of the mentioned

[View Answer](#)

Answer: d

Explanation: Two techniques are present.

4. In crabbing protocol locking

- a) Goes down the tree and back up
- b) Goes up the tree and back down
- c) Goes down the tree and releases
- d) Goes up the tree and releases

[View Answer](#)

Answer: a

Explanation: It moves in a crab like manner.

5. The deadlock can be handled by

- a) Removing the nodes that are deadlocked
- b) Restarting the search after releasing the lock
- c) Restarting the search without releasing the lock
- d) Resuming the search

[View Answer](#)

Answer: b

Explanation: Crabbing protocol moves in a crab like manner.

6. In crabbing protocol, the lock obtained on the root node is in _____ mode.

- a) Shared
- b) Exclusive

- c) Read only
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: Crabbing protocol moves in a crab like manner down the index tree.

7. If needed to split a node or coalesce it with its siblings, or redistribute key values between siblings, the crabbing protocol locks the parent of the node in _____ mode.

- a) Shared
- b) Exclusive
- c) Read only
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Crabbing protocol moves in a crab like manner down the index tree.

8. In crabbing protocol to inset or delete a key value the leaf node has to be locked in _____ mode.

- a) Shared
- b) Exclusive
- c) Read only
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Crabbing protocol moves in a crab like manner down the index tree.

9. B-link tree requires a pointer to its _____ sibling.

- a) Upper
- b) Lower
- c) Right
- d) Left

[View Answer](#)

Answer: c

Explanation: This pointer is required because a lookup that occurs while a node is being split may have to search not only that node but also that node's right sibling.

10. Instead of locking index leaf nodes in a two-phase manner, some index concurrency-control schemes use _____ on individual key values, allowing other key values to be inserted or deleted from the same leaf.

- a) B+ tree locking
- b) Link level locking
- c) Key-value locking
- d) Next value locking

[View Answer](#)

Answer: c

Explanation: Key-value locking thus provides increased concurrency.

Database Questions and Answers – Failure Classification

1. The recovery scheme must also provide

- a) High availability
- b) Low availability
- c) High reliability
- d) High durability

[View Answer](#)

Answer: a

Explanation: It must minimize the time for which the database is not usable after a failure.

2. Which one of the following is a failure to a system

- a) Boot crash
- b) Read failure
- c) Transaction failure
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Types of system failure are transaction failure, system crash and disk failure.

3. Which of the following belongs to transaction failure

- a) Read error
- b) Boot error
- c) Logical error
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Types of system transaction failure are logical and system error.

4. The system has entered an undesirable state (for example, deadlock), as a result of which a transaction cannot continue with its normal execution. This is

- a) Read error
- b) Boot error
- c) Logical error
- d) System error

[View Answer](#)

Answer: c

Explanation: The transaction, can be re-executed at a later time.

5. The transaction can no longer continue with its normal execution because of some internal condition, such as bad input, data not found, overflow, or resource limit exceeded. This is

- a) Read error
- b) Boot error
- c) Logical error
- d) System error

[View Answer](#)

Answer: c

Explanation: The transaction, can be re-executed at a later time.

6. The assumption that hardware errors and bugs in the software bring the system to a halt, but do not corrupt the nonvolatile storage contents,

is known as the

- a) Stop assumption
- b) Fail assumption
- c) Halt assumption
- d) Fail-stop assumption

[View Answer](#)

Answer: d

Explanation: Well-designed systems have numerous internal checks, at the hardware and the software level, that bring the system to a halt when there is an error. Hence, the fail-stop assumption is a reasonable one.

7. Which kind of failure loses its data in head crash or failure during a transfer operation.

- a) Transaction failure
- b) System crash
- c) Disk failure
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Copies of the data on other disks, or archival backups on tertiary media, such as DVD or tapes, are used to recover from the failure.

8. The failure occurred sufficiently early during the transfer that the destination block remains intact.

- a) Partial Failure
- b) Total failure
- c) Successful completion
- d) Data transfer failure

[View Answer](#)

Answer: a

Explanation: Copies of the data on other disks, or archival backups on tertiary media, such as DVD or tapes, are used to recover from the failure.

9. The database is partitioned into fixed-length storage units called

- a) Parts
- b) Blocks
- c) Reads
- d) Build

[View Answer](#)

Answer: b

Explanation: Blocks are the units of data transfer to and from disk, and may contain several data items.

10. Which of the following causes system to crash

- a) Bug in software
- b) Loss of volatile data
- c) Hardware malfunction
- d) All of the mentioned

[View Answer](#)

Answer: d

Explanation: The content of non-volatile storage remains intact, and is not corrupted.

Database Questions and Answers – Recovery

1. The log is a sequence of _____ recording all the update activities in the database.

- a) Log records
- b) Records
- c) Entries
- d) Redo

[View Answer](#)

Answer: a

Explanation: The most widely used structure for recording database modifications is the log.

2. In the _____ scheme, a transaction that wants to update the database first creates a complete copy of the database.

- a) Shadow copy
- b) Shadow Paging
- c) Update log records
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: If at any point the transaction has to be aborted, the system merely deletes the new copy. The old copy of the database has not been affected.

3. The _____ scheme uses a page table containing pointers to all pages; the page table itself and all updated pages are copied to a new location.

- a) Shadow copy
- b) Shadow Paging
- c) Update log records
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: Any page which is not updated by a transaction is not copied, but instead the new page table just stores a pointer to the original page.

4. The current copy of the database is identified by a pointer, called _____ which is stored on disk.

- a) Db-pointer
- b) Update log
- c) Update log records
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Any page which is not updated by a transaction is not copied, but instead the new page table just stores a pointer to the original page.

5. If a transaction does not modify the database until it has committed, it is said to use the _____ technique.

- a) Deferred-modification
- b) Late-modification
- c) Immediate-modification
- d) Undo

[View Answer](#)

Answer: a

Explanation: Deferred modification has the overhead that transactions need to make local copies of all updated data items; further, if a

transaction reads a data item that it has updated, it must read the value from its local copy.

6. If database modifications occur while the transaction is still active, the transaction is said to use the _____ technique.

- a) Deferred-modification
- b) Late-modification
- c) Immediate-modification
- d) Undo

[View Answer](#)

Answer: c

Explanation: We say a transaction modifies the database if it performs an update on a disk buffer, or on the disk itself; updates to the private part of main memory do not count as database modifications.

7. _____ using a log record sets the data item specified in the log record to the old value.

- a) Deferred-modification
- b) Late-modification
- c) Immediate-modification
- d) Undo

[View Answer](#)

Answer: d

Explanation: Undo brings the previous contents.

8. In the _____ phase, the system replays updates of all transactions by scanning the log forward from the last checkpoint.

- a) Repeating
- b) Redo
- c) Replay
- d) Undo

[View Answer](#)

Answer: b

Explanation: Undo brings the previous contents.

9. The actions which are played in the order while recording it is called _____ history.

- a) Repeating
- b) Redo
- c) Replay
- d) Undo

[View Answer](#)

Answer: a

Explanation: Undo brings the previous contents.

10. A special redo-only log record $\langle T_i, X_j, V_1 \rangle$ is written to the log, where V_1 is the value being restored to data item X_j during the rollback. These log records are sometimes called

- a) Log records
- b) Records
- c) Compensation log records
- d) Compensation redo records

[View Answer](#)

Answer: c

Explanation: Such records do not need undo information since we never need to undo such an undo operation.

Database Questions and Answers – Buffer Management

1. In order to reduce the overhead in retrieving the records from the storage space we use

- a) Logs
- b) Log buffer
- c) Medieval space
- d) Lower records

[View Answer](#)

Answer: b

Explanation: The output to stable storage is in units of blocks.

2. The order of log records in the stable storage _____ as the order in which they were written to the log buffer.

- a) Must be exactly the same
- b) Can be different
- c) Is opposite
- d) Can be partially same

[View Answer](#)

Answer: a

Explanation: As a result of log buffering, a log record may reside in only main memory (volatile storage) for a considerable time before it is output to stable storage.

3. Before a block of data in main memory can be output to the database, all log records pertaining to data in that block must have been output to stable storage. This is

- a) Read-write logging
- b) Read-ahead logging
- c) Write-ahead logging
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: The WAL rule requires only that the undo information in the log has been output to stable storage, and it permits the redo information to be written later.

4. Writing the buffered log to _____ is sometimes referred to as a log force.

- a) Memory
- b) Backup
- c) Redo memory
- d) Disk

[View Answer](#)

Answer: d

Explanation: If there are insufficient log records to fill the block, all log records in main memory are combined into a partially full block and are output to stable storage.

5. The _____ policy, allows a transaction to commit even if it has modified some blocks that have not yet been written back to disk.

- a) Force
- b) No-force
- c) Steal
- d) No-steal

[View Answer](#)

Answer: b

Explanation: No-force policy allows faster commit of transactions.

6. _____ policy allows multiple updates to accumulate on a block before it is output to stable storage, which can reduce the number of output operations greatly for frequently updated blocks.

- a) Force
- b) No-force
- c) Steal
- d) No-steal

[View Answer](#)

Answer: b

Explanation: No-force policy allows faster commit of transactions.

7. The _____ policy, allows the system to write modified blocks to disk even if the transactions that made those modifications have not all committed.

- a) Force
- b) No-force
- c) Steal
- d) No-steal

[View Answer](#)

Answer: c

Explanation: The no-steal policy does not work with transactions that perform a large number of updates.

8. Locks on buffer blocks are unrelated to locks used for concurrency-control of transactions, and releasing them in a non-two-phase manner does not have any implications on transaction serializability. This is

- a) Latches
- b) Swap Space
- c) Dirty Block
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: These locks, and other similar locks that are held for a short duration.

9. The _____ contains a list of blocks that have been updated in the database buffer.

- a) Latches
- b) Swap Space
- c) Dirty Block
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Dirty blocks are those that have been updated in memory, and the disk version is not up-to-date.

10. The operating system reserves space on disk for storing virtual-memory pages that are not currently in main memory; this space is called

- a) Latches
- b) Swap Space
- c) Dirty Block
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: Almost all current-generation operating systems retain complete control of virtual memory.

Database Questions and Answers – Failure with Nonvolatile Storage

1. The silicon chips used for data processing are called

- a) RAM chips
- b) ROM chips
- c) Micro processors
- d) PROM chips

[View Answer](#)

Answer: d

Explanation: PROM is Programmable Read Only Memory.

2. Which of the following is used for manufacturing chips?

- a) Control bus
- b) Control unit
- c) Parity unit
- d) Semiconductor

[View Answer](#)

Answer: d

Explanation: A semiconductor is a material which has electrical conductivity between that of a conductor such as copper and that of an insulator such as glass.

3. What was the name of the first commercially available microprocessor chip?

- a) Intel 308
- b) Intel 33
- c) Intel 4004
- d) Motorola 639

[View Answer](#)

Answer: c

Explanation: The Intel 4004 is a 4-bit central processing unit (CPU) released by Intel Corporation in 1971

4. The magnetic storage chip used to provide non-volatile direct access storage of data and that have no moving parts are known as

- a) Magnetic core memory
- b) Magnetic tape memory
- c) Magnetic disk memory
- d) Magnetic bubble memory

[View Answer](#)

Answer: d

Explanation: Bubble domain visualization by using CMOS-MagView.

5. The ALU of a computer normally contains a number of high speed storage element called

- a) Semiconductor memory
- b) Registers
- c) Hard disks
- d) Magnetic disk

[View Answer](#)

Answer: b

Explanation: External control unit tells the ALU what operation to perform on that data, and then the ALU stores its result into an output register.

6. Which of the following is used only for data entry and storage, and never for processing?

- a) Mouse
- b) Dumb terminal
- c) Micro computer
- d) Dedicated data entry system

[View Answer](#)

Answer: b

Explanation: Dumb terminals are those that can interpret a limited number of control codes.

7. Non-volatile storage needs to have a _____ where the loses in future can be recovered.

- a) Dump
- b) Recover place
- c) Disk
- d) Redo plan

[View Answer](#)

Answer: a

Explanation: The basic scheme is to dump the entire contents of the database to stable storage periodically—say, once per day.

8. A dump of the database contents is also referred to as an _____ dump.

- a) Archival
- b) Fuzzy
- c) SQL
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: We can archive the dumps and use them later to examine old states of the database.

9. _____ dump, writes out SQL DDL statements and SQL insert statements to a file, which can then be reexecuted to re-create the database.

- a) Archival
- b) Fuzzy
- c) SQL
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: Such dumps are useful when migrating data to a different instance of the database, or to a different version of the database software, since the physical locations and layout may be different in the other database instance or database software version.

10. _____ dump schemes have been developed that allow transactions to be active while the dump is in progress.

- a) Archival
- b) Fuzzy
- c) SQL
- d) All of the mentioned

[View Answer](#)

Answer: b

Explanation: The simple dump procedure described here is costly and so fuzzy dump is used.

Database Questions and Answers – ARIES

1. ARIES uses a _____ to identify log records, and stores it in database pages.

- a) Log sequence number
- b) Log number
- c) Lock number
- d) Sequence

[View Answer](#)

Answer: b

Explanation: LSN is used to identify which operations have been applied to a database page.

2. ARIES supports _____ operations, which are physical in that the affected page is physically identified, but can be logical within the page.

- a) Physiological redo
- b) Physiological undo
- c) Logical redo
- d) Logical undo

[View Answer](#)

Answer: a

Explanation: The deletion of a record from a page may result in many other records in the page being shifted, if a slotted page structure is used.

3. _____ is used to minimize unnecessary redos during recovery.

- a) Dirty page table
- b) Page table
- c) Dirty redo
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Dirty pages are those that have been updated in memory, and the disk version is not up-to-date.

4. _____ scheme that records only information about dirty pages and associated information and does not even require of writing dirty pages to disk.

- a) Fuzzy logic
- b) Checkpoints
- c) Fuzzy-checkpoint
- d) Logical checkpoint

[View Answer](#)

Answer: c

Explanation: It flushes dirty pages in the background, continuously, instead of writing them during checkpoints.

5. Whenever an update operation occurs on a page, the operation stores the LSN of its log record in the _____ field of the page.

- a) LSN
- b) ReadLSN
- c) PageLSN
- d) RedoLSN

[View Answer](#)

Answer: c

Explanation: Each page maintains an identifier called the PageLSN.

6. There are special redo-only log records generated during transaction rollback, called _____ in ARIES.

- a) Compensation log records
- b) Read log records
- c) Page log records
- d) Redo log records

[View Answer](#)

Answer: a

Explanation: These serve the same purpose as the redo-only log records in our earlier recovery scheme.

7. The _____ contains a list of pages that have been updated in the database buffer.

- a) Dirty page table
- b) Page table
- c) Dirty redo
- d) All of the mentioned

[View Answer](#)

Answer: a

Explanation: Dirty pages are those that have been updated in memory, and the disk version is not up-to-date.

8. _____ determines which transactions to undo, which pages were dirty at the time of the crash, and the LSN from which the redo pass should start.

- a) Analysis pass
- b) Redo pass
- c) Undo pass
- d) None of the mentioned

[View Answer](#)

Answer: a

Explanation: The analysis pass finds the last complete checkpoint log record, and reads in the DirtyPageTable from this record.

9. _____ starts from a position determined during analysis, and performs a redo, repeating history, to bring the database to a state it was in before the crash.

- a) Analysis pass
- b) Redo pass
- c) Undo pass
- d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: The redo pass repeats history by replaying every action that is not already reflected in the page on disk.

10. _____ rolls back all transactions that were incomplete at the time of crash.

- a) Analysis pass
- b) Redo pass
- c) Undo pass
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: It performs a single backward scan of the log, undoing all transactions in undo-list.

Database Questions and Answers – Lock Release and Undo Operations

1. Which lock should be obtained to prevent a concurrent transaction from executing a conflicting read, insert or delete operation on the same key value.

- a) Higher-level lock
- b) Lower-level lock
- c) Read only lock
- d) Read write

[View Answer](#)

Answer: a

Explanation: Operations acquire lower-level locks while they execute, but release them when they complete; the corresponding transaction must however retain a higher-level lock in a two-phase manner to prevent concurrent transactions from executing conflicting actions.

2. Once the lower-level lock is released, the operation cannot be undone by using the old values of updated data items, and must instead be undone by executing a compensating operation; such an operation is called

- a) Logical operation
- b) Redo operation
- c) Logical undo operation
- d) Undo operation

[View Answer](#)

Answer: a

Explanation: It is important that the lower-level locks acquired during an operation are sufficient to perform a subsequent logical undo of the operation.

3. Which of the following is used for undo operations alone?

- a) Logical logging
- b) Physical logging
- c) Physical log records
- d) Physical logging and Physical log records

[View Answer](#)

Answer: a

Explanation: If the operation inserted an entry in a B+-tree, the undo information U would indicate that a deletion operation is to be performed, and would identify the B+-tree and what entry to delete from the tree. Such logging of information about operations is called logical logging.

4. Redo operations are performed exclusively using

- a) Logical logging
- b) Physical logging
- c) Physical log records
- d) Both Physical logging and Physical log records

[View Answer](#)

Answer: d

Explanation: Logging of old-value and new-value information is called physical logging.

5. To perform logical redo or undo, the database state on disk must be operation _____ that is, it should not have partial effects of any operation.

- a) Persistent
- b) Resistant
- c) Consistent
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: Data structures such as B+-trees would not be in a consistent state, and neither logical redo nor logical undo operations can be performed on an inconsistent data structure.

6. An operation is said to be _____ if executing it several times in a row gives the same result as executing it once.

- a) Idempotent
- b) Changed
- c) Repetitive
- d) All of the above

[View Answer](#)

Answer: a

Explanation: Operations such as inserting an entry into a B+-tree may not be idempotent, and the recovery algorithm must therefore make sure that an operation that has already been performed is not performed again.

7. Immediate database modification technique uses

- a) Both undo and redo
- b) Undo but no redo
- c) Redo but no undo
- d) Neither undo nor redo

[View Answer](#)

Answer: a

Explanation: Undo erases all the changes and redo makes the deleted changes.

8. Shadow paging has

- a) no redo
- b) no undo
- c) redo but no undo
- d) neither redo nor undo

[View Answer](#)

Answer: a

Explanation: Undo erases all the changes and redo makes the deleted changes.

9. For correct behaviour during recovery, undo and redo operation must be

- a) Commutative
- b) Associative
- c) Idempotent
- d) Distributive

[View Answer](#)

Answer: c

Explanation: Undo erases all the changes and redo makes the deleted changes.

10. If _____ are not obtained in undo operation it will cause problem in undo-phase.

- a) Higher-level lock
- b) Lower-level lock
- c) Read only lock
- d) Read write

[View Answer](#)

Answer: b

Explanation: Operations acquire lower-level locks while they execute, but release them when they complete; the corresponding transaction must however retain a higher-level lock in a two-phase manner to prevent concurrent transactions from executing conflicting actions.

Database Questions and Answers – Remote Backup Systems

1. The remote backup site is sometimes also called the
 - a) Primary Site
 - b) Secondary Site
 - c) Tertiary Site
 - d) None of the mentioned

[View Answer](#)

Answer: b

Explanation: We can achieve high availability by performing transaction processing at one site, called the primary site, and having a remote backup site where all the data from the primary site are replicated.

2. Remote backup system must be _____ with the primary site.

- a) Synchronised
- b) Separated
- c) Connected
- d) Detached but related

[View Answer](#)

Answer: a

Explanation: We can achieve high availability by performing transaction processing at one site, called the primary site, and having a remote backup site where all the data from the primary site are replicated.

3. The backup is taken by
 - a) Erasing all previous records
 - b) Entering the new records
 - c) Sending all log records from primary site to the remote backup site
 - d) Sending selected records from primary site to the remote backup site

[View Answer](#)

Answer: c

Explanation: We can achieve high availability by performing transaction processing at one site, called the primary site, and having a remote backup site where all the data from the primary site are replicated.

4. When the _____ the backup site takes over processing and becomes the primary.

- a) Secondary fails
- b) Backup recovers
- c) Primary fails
- d) None of the mentioned

[View Answer](#)

Answer: c

Explanation: When the original primary site recovers, it can either play the role of remote backup, or take over the role of primary site again.

5. The simplest way of transferring control is for the old primary to receive _____ from the old backup site.

- a) Undo logs
- b) Redo Logs
- c) Primary Logs
- d) All of the mentioned

[View Answer](#)

Answer: c

Explanation: If control must be transferred back, the old backup site can pretend to have failed, resulting in the old primary taking over.

6. The time to process the remote backup can be reduced by

- a) Flags
- b) Breakpoints
- c) Redo points
- d) Checkpoints

[View Answer](#)

Answer: d

Explanation: If the log at the remote backup grows large, recovery will take a long time. The remote backup site can periodically process the redo log records that it has received and can perform a checkpoint, so that earlier parts of the log can be deleted.

7. A _____ configuration can make takeover by the backup site almost instantaneous.

- a) Hot-spare
- b) Remote
- c) Direct
- d) Spare

[View Answer](#)

Answer: d

Explanation: In this configuration, the remote backup site continually processes redo log records as they arrive, applying the updates locally.

8. A transaction commits as soon as its commit log record is written to stable storage at the primary site. This is

- a) One Safe
- b) Two Safe
- c) Two-very Safe
- d) Very Safe

[View Answer](#)

Answer: a

Explanation: The problem with this scheme is that the updates of a committed transaction may not have made it to the backup site, when the backup site takes over processing.

9. A transaction commits as soon as its commit log record is written to stable storage at the primary and the backup site. This is

- a) One Safe
- b) Two Safe
- c) Two-very Safe
- d) Very Safe

[View Answer](#)

Answer: c

Explanation: The problem with this scheme is that transaction processing cannot proceed if either the primary or the backup site is down.

10. If only the primary is active, the transaction is allowed to commit as soon as its commit log record is written to stable storage at the primary site. This is

- a) One Safe
- b) Two Safe
- c) Two-very Safe
- d) Very Safe

[View Answer](#)

Answer: b

Explanation: This scheme provides better availability than does two-very-safe, while avoiding the problem of lost transactions faced by the one-safe scheme.